

Mechanism Design and Risk Aversion

Anand Bhalgat Tanmoy Chakraborty Sanjeev Khanna
Univ. of Pennsylvania Harvard University Univ. of Pennsylvania
bhalgat@seas.upenn.edu tanmoy@seas.harvard.edu sanjeev@cis.upenn.edu

Abstract

We develop efficient algorithms to construct utility maximizing mechanisms in the presence of risk averse players (buyers and sellers) in Bayesian single parameter and multi-parameter settings. We model risk aversion by a concave utility function, and players play strategically to maximize their expected utility. Bayesian mechanism design has usually focused on maximizing expected revenue in a *risk neutral* environment, *i.e.* where all the buyers and the seller have linear utility, and no succinct characterization of expected utility maximizing mechanisms is known even for single-parameter multi-unit auctions.

We first consider the problem of designing optimal DSIC (dominant strategy incentive compatible) mechanism for a risk averse seller in the case of multi-unit auctions, and we give a poly-time computable deterministic sequential posted pricing mechanism (SPM) that for any $\epsilon > 0$, yields a $(1 - 1/e - \epsilon)$ -approximation to the expected utility of the seller in an optimal DSIC mechanism. Our result is based on a novel application of a correlation gap bound, along with *splitting* and *merging* of random variables to redistribute probability mass across buyers. This allows us to reduce our problem to that of checking feasibility of a small number of distinct configurations, each of which corresponds to a covering LP. A feasible solution to the LP gives us the distribution on prices for each buyer to use in a randomized SPM. We get a deterministic SPM by sampling from this randomized SPM. Our techniques extend to the multi-parameter setting with unit demand buyers.

We next consider the setting when buyers as well as the seller are risk averse, and the objective is to maximize the seller's expected utility. We design a truthful-in-expectation mechanism whose utility is a $\left(\left(1 - \frac{1}{e}\right)^2 \times \max \left(1 - \frac{1}{e}, 1 - \frac{1}{\sqrt{2\pi k}}\right) \right)$ -approximation to the optimal BIC mechanism under two mild assumptions: (a) ex post individual rationality and (b) no positive transfers. Our mechanism consists of multiple rounds. It considers each buyer in a round with small probability, and when a buyer is considered, it allocates an item to the buyer according to payment functions that are computed using stochastic techniques developed for DSIC mechanisms. Lastly, we consider the problem of revenue maximization for a risk neutral seller in presence of risk averse buyers, and give a poly-time algorithm to design an optimal mechanism for the seller.

We believe that the techniques developed in this work will be useful in handling other stochastic optimization problems with a concave objective function.

1 Introduction

Bayesian mechanism design has usually focused on maximizing expected revenue in a *risk neutral* environment, *i.e.* where all the buyers and the seller have linear utility, and choose their strategy with the aim of maximizing their *expected payoff*. However, since the payoff is a random outcome that depends on other players' valuations and strategies, there is risk associated with it. A standard model [4, 13] that captures risk aversion assumes that a player has a non-decreasing concave utility function $U : (-\infty, \infty) \rightarrow (-\infty, \infty)$, so that when the payoff obtained is R , the player's utility is $U(R)$. The player may choose to express various levels of risk aversion by specifying a suitable concave function as his utility, and then his aim becomes to maximize his expected utility. While mechanism design in a risk neutral (linear utility) environment is well understood for multi-unit auctions, many properties tend to break down in the presence of risk aversion (concave utility). In this paper, we develop efficient algorithms to compute mechanisms in the presence of risk averse players. We mainly focus on the prominent single parameter setting of multi-unit auctions.

Risk neutral Seller, Risk Averse Buyers: Let us first analyze the effect of risk aversion among buyers when the seller is risk neutral, *i.e.* he wants to maximize his expected revenue. This has been the predominant model for studying risk aversion in mechanism design. Myerson's characterization [12] of the optimal auction design does not apply when buyers are risk averse. In particular, revenue equivalence [12] does not hold, and an optimal dominant strategy incentive compatible (DSIC) mechanism may generate less expected revenue than an optimal Bayesian incentive compatible (BIC) mechanism. This is because when buyers are risk averse, the seller can extract greater expected revenue by offering a deterministic payment scheme to the buyers and charging extra for this *insurance*. As a specific example, consider buyers with *constant absolute risk aversion*, *i.e.* exponentially diminishing marginal utility. For any DSIC mechanism (such as Myerson's), one can construct a corresponding BIC mechanism with the same allocation curves: if a buyer i reports valuation v , charge a deterministic payment $\alpha(i, v)$ (perhaps zero) if the buyer does not get the item, and $\beta(i, v)$ if he gets the item. For the latter mechanism to be truthful, it must satisfy *utility equivalence* [10, 11, 8]: the net expected utility of the buyer in the latter mechanism must match the former when his value is v , for each v . However, if his expected payment in the latter mechanism remains equal to the former, then his net expected utility will go up, because concave utility makes the deterministic payment more preferable. Thus the expected payment must be higher in the latter mechanism, implying greater revenue. Note that the direct revelation principle still holds, so there is a BIC mechanism that generates as much expected revenue as any Bayes-Nash equilibrium^a.

Maskin and Riley [9] characterized optimal BIC mechanism in this setting for selling a single item when buyers' value distributions are IID. It is assumed that the buyers' utility functions are known to the mechanism designer (the seller), since truthfulness itself depends on these functions. Another well-known result states that under some natural assumptions on the buyers' utility functions, first-price auction with reserve (specifically, Bayes-Nash equilibrium of this auction) generates greater revenue than second-price auction with the same reserve [10, 9, 11]. These results show that the presence of risk averse buyers should have substantial effect on optimal mechanism design.

Our results for a k -unit auction are summarized in Table 1. We first design a poly-time algorithm to compute a BIC mechanism for a risk neutral seller when buyers are risk averse with publicly known utility functions. This result extends the work of Maskin and Riley [9], in a computational sense, to the general setting of multi-unit auctions with non-identical distributions. Our algorithm is

^aAny equilibrium of the mechanism where buyers play strategically to maximize their utility.

Type of risk environment	Comparison with Optimal DSIC	Comparison with Optimal BIC	
	Poly-time DSIC ^a	Poly-time TIE ^b	Poly-time BIC ^b
Risk neutral seller, risk neutral buyers	1 [Myerson '81]	1 [Myerson '81]	1 [Myerson '81]
Risk neutral seller, risk averse buyers	1 [Myerson '81]	$\gamma(k)$ ^c [Theorem 4]	1 [Theorem 4]
Risk averse seller, risk neutral buyers	$(1 - 1/e - \epsilon)$ [Theorem 1]	$(1 - 1/e)^2 \gamma(k) - \epsilon$ [Theorem 2] ^d	1 [Eso-Futo '99]
Risk averse seller, risk averse buyers	$(1 - 1/e - \epsilon)$ [Theorem 1]	$(1 - 1/e)^2 \gamma(k) - \epsilon$ [Theorem 2] ^d	$(1 - 1/e)^2 \gamma(k) - \epsilon$ [Theorem 2] ^d

^aNeed to know seller's utility function. Independent of buyers' utility functions as long as they are non-decreasing.

^bNeed to know both seller and buyers' utility functions.

^c $\gamma(k) = (1 - \frac{k^k}{k!e^k})$. $\gamma(1) = 1 - 1/e$, and it approaches $(1 - \frac{1}{\sqrt{2\pi k}})$ for large k .

^dImproves to $(1 - 1/e)\gamma(k) - \epsilon$ for IID buyers (Theorem 3). Further, the factor improves to $(1 - 1/e - \epsilon)$ if $k \geq 1/\epsilon^3$. Here, comparison is made only against optimal BIC satisfying: (i) ex-post IR, and (ii) no positive transfers.

Table 1: Summary of approximation results for k -unit auctions.

a linear program developed using a general form of Border's inequality. Further, we design a a poly-time computable randomized truthful-in-expectation mechanism which is a $\gamma(k)$ -approximation to the utility optimal BIC mechanism. Here, $\gamma(k) = (1 - \frac{k^k}{k!e^k})$; $\gamma(1) = 1 - 1/e$, and approaches $(1 - \frac{1}{\sqrt{2\pi k}})$ for large k . A mechanism is said to be *truthful in expectation* (TIE) if truth-telling maximizes each buyer's expected utility, regardless of other buyers' bids (expected utility is measured only over the random bits used by the mechanism). This is a stronger truthfulness requirement than BIC, but weaker than DSIC ^b, and does not rely on all buyers sharing the same belief about each other. TIE mechanisms can still generate greater expected revenue than DSIC mechanisms. In effect, this result *bounds the gap between TIE mechanisms and BIC mechanisms*.

Risk Averse Seller, Risk Averse Buyers: Next, we consider the scenario where the seller as well as the buyers are risk averse. Risk aversion from a seller's perspective has received relatively less attention, and no work has considered both sides to be risk averse. Eso and Futo [7] designed an optimal BIC mechanism for a risk averse seller when *buyers are risk neutral*. In this setting, the seller can *transfer the entire risk* to risk neutral buyers, and obtain the expected revenue of Myerson's mechanism in every realization. A DSIC mechanism cannot do such a risk transfer, and as such, the gap between optimal BIC and optimal DSIC is unbounded, as illustrated below. *Example.* Consider an instance with two buyers and two items. Each buyer has valuation 1 for the item w.p. ϵ and 0 otherwise. The seller's utility function \mathbf{U} is as follows: $\mathbf{U}(t) = \min\{t, \epsilon\}$. The utility optimal DSIC mechanism sets a price of 1 to each buyer, and gets utility ϵ with probability $2\epsilon(1 - \epsilon) \leq 2\epsilon$, otherwise its utility is 0. So the expected utility of an optimal DSIC mechanism is at most $2\epsilon^2$. If the first buyer is risk neutral, then we can design a BIC mechanism as follows: charge the first buyer ϵ in every realization (even when his value is zero), and set a price of 1 to the second buyer. If the second buyer pays up 1 (which happens w.p. ϵ), then pay 1 dollar to the first buyer. The first buyer never gets the item, and the mechanism is incentive compatible for the first

^bIn this paper, we require that in a DSIC mechanism, truth-telling must maximize buyers' utility regardless of the mechanism's random bits. This distinction between DSIC and TIE will play a crucial role below.

buyer. The seller gets a revenue of ϵ in every realization, so his expected utility is ϵ . Therefore the gap is unbounded as $\epsilon \rightarrow 0$.

However, the extent to which the seller can transfer its risk to the buyers depends upon buyers' utility functions. So the result of [7] does not hold when buyers are risk averse. We design TIE mechanisms that are constant approximation to a utility-optimal BIC mechanism for a risk averse seller. We restrict our comparison only to BIC mechanisms that do not allow *positive transfers* (i.e. there is no payment from the seller to any buyer in any realization), and are *ex post individually rational*. Our approximation factor is $(1 - 1/e)^2 \gamma(k)$ ($(1 - 1/e) \gamma(k)$ for IID buyers), and approaches $(1 - 1/e)$ (using a slightly different algorithm) as k becomes large. This implies a constant upper bound on the gap between TIE mechanisms and BIC mechanisms without positive transfer.

DSIC mechanisms for Risk Averse Sellers: So far, we have designed BIC or TIE mechanisms assuming that the players' utility functions are known to the designer. Though this has been a standard assumption through most of the literature on risk aversion, it is usually the seller who designs the mechanism, and it is not practical to assume that the seller knows all the buyers' utility functions. Unfortunately, it is impossible in general to even check if a mechanism is BIC or even TIE (or identify a Bayes-Nash equilibrium of a given mechanism) without knowing the buyers' utility functions. In contrast, *DSIC mechanisms are independent of buyers' utility functions* as long as buyers' utility functions are non-decreasing. So we do not need to know the buyers' utility functions to compute a utility-optimal DSIC mechanism for the seller.

Clearly, for a risk-neutral seller, Myerson's mechanism remains an optimal DSIC mechanism even when buyers are risk averse. However, that is not true for a risk averse seller. Further, a virtual value maximization approach does not apply when the seller has at least two units of inventory^c. This is because the contributions of different buyers cannot be counted separately when the seller has non-linear utility. We design a poly-time computable $(1 - 1/e - \epsilon)$ -approximate DSIC mechanism. Our mechanism is a sequential posted-pricing mechanism (SPM), that decides a price for each buyer, and then makes take-it-or-leave-it offers to the buyers in decreasing order of prices, till inventory runs out. Posted pricing mechanisms have been studied extensively for maximizing expected revenue (eg. [6, 5, 16]).

Sundarajan and Yan [14] designed DSIC mechanisms for multi-unit auctions for a risk-averse seller when buyers' valuation functions are *regular*. They focused only on mechanisms that do not depend even on the seller's own utility function. While this may be an attractive property, the optimal mechanism is no longer well-defined in this case (even neglecting computational limitations), and it forces the approximation guarantees to be weaker – $1/8$ -approximation for regular distributions ($1/2$ when there is unlimited supply of items), and there is a lower bound instance implying unbounded gap for general distributions. Non-regular distributions are not uncommon – any distribution with more than one mode is non-regular. Risk aversion is particularly important in the presence of such high variance distributions, and it is a reasonable approach for a seller to decide upon his own utility function and then use our algorithm to design a mechanism.

As a final note, we are able to extend our techniques to give a constant approximation to an optimal deterministic DSIC mechanism in a multi-parameter setting, namely, when there are multiple distinct items and unit-demand buyers (see Appendix E).

^cIf there is only one item to sale, then at most one buyer pays in any realization, and the seller's utility can be maximized by scaling the bid values using the utility function.

1.1 Overview of Techniques

We first design a DSIC mechanism for a risk averse seller, which turns out to be a relatively simpler problem than competing against optimal BIC mechanism. We establish our main probabilistic tools in the process, which later get used for our BIC results. For DSIC mechanisms, we first argue that a $(1 - \frac{1}{e})$ -approximate (randomized) SPM can be obtained by using the same price distribution as that offered to each buyer in the optimal mechanism, except that the prices are now set independently (see Lemma 3.3). The argument uses the *correlation gap* bound of Agrawal et. al. [1] for submodular objectives. However, this is only an existential result, since getting the SPM requires oracle access to a utility-optimal DSIC mechanism.

Our main technical contribution is to show that it suffices (with same loss factor of $(1 - 1/e)$) to match the optimal mechanism only in the *sum of sale probabilities over all buyers, and not the sale probability for each buyer*, at every price. That is, any two mechanisms that match in this *coarse footprint* will have approximately equal expected utility. This property follows from a generalization of the correlation gap bound in [1], which not only introduces independence but also redistributes probability mass across variables (see Lemma 3.7). The redistribution is achieved by *splitting* and *merging* random variables to transform one given mechanism to another that matches the coarse footprint. Using a sophisticated classification of prices, we show that it suffices to match an even coarser footprint containing only constant number of parameters, which define a *configuration*. The algorithm finds a feasible solution for each configuration using a *covering LP*. Then, it simulates these SPMs, one for each feasible configuration, to choose one with the highest expected utility.

To design a BIC mechanism when the seller as well as the buyers are risk averse, the techniques developed for DSIC mechanisms can be used to establish that if allocation and payment functions of the optimal mechanism across buyers are made independent, and inventory constraints removed, the utility will be at least $(1 - 1/e)\text{OPT}_{\text{BIC}}$. However, to convert such a *soft* mechanism into a mechanism that strictly satisfies the inventory constraint is not easy: if we restrict the allocation to buyers with *top k* payments in a realization of a soft mechanism, a function which is submodular, the resulting mechanism is no longer BIC. Further, distributions on the revenue from any two allocations in the mechanism are incomparable, so restricting to *first k* allocations in a realization of a soft mechanism can be arbitrarily bad. To overcome this problem, we develop a mechanism with $L \rightarrow \infty$ *rounds*, such that in each round, each buyer is ignored with a high probability of $(1 - 1/L)$. We show that the revenue from each allocation in this mechanism has identical distribution, and the loss in the expected utility caused by imposing the hard inventory constraint is bounded.

Organization: Rest of the paper is organized as follows. In Section 2, we provide some background material. In Sections 3 and 4, we present our DSIC and BIC mechanisms for risk averse seller, respectively. Our BIC mechanisms for the risk neutral seller and risk averse buyers appear in Appendix D, and our result for multi-parameter unit-demand setting appears in Appendix E.

2 Preliminaries

Single Parameter Multi-Unit Auctions: The seller provides a single type of item (or service), of which he has k identical copies. There are n buyers $\{1, 2, \dots, n\}$, who have some private value for that service. Let buyer i have a valuation of v_i for the item (and he can consume only one unit), which is drawn, independent of other buyers' valuations, from a known distribution with cdf $F_i(x) = \Pr[v_i \leq x]$. We refer to $\mathbf{v} = (\mathbf{v}_1, \mathbf{v}_2 \dots \mathbf{v}_n)$ as the *valuation vector*.

Revenue, Utility and Optimality: The *revenue* $\text{Rev}(\mathcal{M}, \mathbf{v})$ of a mechanism \mathcal{M} , when the realized valuation vector is \mathbf{v} , is the sum of payments from each buyers. The expected revenue of a mechanism \mathcal{M} is $\mathbf{E}_{\mathbf{v}}[\text{Rev}(\mathcal{M}, \mathbf{v})]$. In this work, we assume that the seller has a monotonically increasing concave utility function \mathbf{U} , which also satisfies $\mathbf{U}(0) = 0$. The utility of the mechanism is $\mathbf{U}(\text{Rev}(\mathcal{M}, \mathbf{v}))$, and the expected utility of the mechanism is $\mathbf{U}(\mathcal{M}) = \mathbf{E}_{\mathbf{v}}[\mathbf{U}(\text{Rev}(\mathcal{M}, \mathbf{v}))]$. Let OPT_{DSIC} and OPT_{BIC} denote the expected utility of a utility-optimal DSIC and BIC mechanisms respectively. A mechanism is said to be an α -approximation to optimal DSIC (or BIC) mechanism if $\mathbf{U}(\mathcal{M}) \geq \alpha \text{OPT}_{\text{DSIC}}$ (or $\mathbf{U}(\mathcal{M}) \geq \alpha \text{OPT}_{\text{BIC}}$).

DSIC Mechanisms: It is well-known (eg. [12]) that a DSIC mechanism sets a (possibly randomized) price for buyer i based on v_{-i} but independent of v_i , and buyer i gets an item if and only if his valuation exceeds this price. Given this characterization, it is easy to see that as long as a buyer has a non-decreasing utility function, he will report truthfully in a DSIC mechanism, for any realization of valuation vector and random bits of the mechanism. Note that random bits do not help a DSIC mechanism obtain greater utility, since the definition of DSIC implies that truthfulness must hold even if the random bits were revealed prior to submitting bids. So there is a utility-optimal DSIC mechanism which is deterministic.

Buyer's Risk Aversion and BIC Mechanisms: Each buyer i is associated with a publicly known monotone concave utility function \mathcal{U}_i (defined on the value of item received minus payment) with $\mathcal{U}_i(0) = 0$. A BIC mechanism is associated with two functions $h(\cdot, \cdot, \cdot)$ and $g(\cdot, \cdot, \cdot)$: $h(i, j, v)$ is the probability that for valuation v , buyer i is allocated an item for a payment of p_j , and $g(i, j, v)$ is the probability that he pays p_j and is not allocated an item for valuation v . We refer to these two functions as the *payment functions* of the mechanism. We note that the allocation and payment of a buyer is possibly correlated with other buyers' payments, allocations as well as their valuations. Thus, a mechanism is BIC if and only if for each i, v, v' , we have

$$\sum_j (\mathcal{U}_i(v - p_j)h(i, j, v) + \mathcal{U}_i(-p_j)g(i, j, v)) \geq \sum_j (\mathcal{U}_i(v - p_j)h(i, j, v') + \mathcal{U}_i(-p_j)g(i, j, v'))$$

We note, given any buyer i , we allow his payment to be randomized rather than a fixed value as a function of buyer i 's valuation and whether he gets an item. This strictly gives more power to a risk averse seller maximizing his expected utility. This is in contrast to the setting considered by Maskin and Riley [9], where it suffices to assume that buyer i 's payment for valuation v is a fixed value as a function of v and whether he gets the item.

We define a *soft randomized sequential mechanism* as a mechanism without inventory limit that arranges buyers in an arbitrary order, asks each buyer for his valuation one-by-one. If the buyer i 's reported valuation is v , the mechanism allocates an item to him independently w.p. $\sum_j h(i, j, v)$. If he is allocated an item, then the seller charges him p_j w.p. $\frac{h(i, j, v)}{\sum_l h(i, l, v)}$. When he is not allocated an item, he pays p_j w.p. $\frac{g(i, j, v)}{\sum_l g(i, l, v)}$. *Randomized sequential mechanisms* are same as *soft randomized sequential mechanisms* with an exception that they stop after running out of inventory. We note that if a *soft randomized sequential mechanism* is BIC, then the corresponding *randomized sequential mechanism* is also BIC.

Stochastic Dominance: Given two non-negative distributions \mathcal{D}_1 and \mathcal{D}_2 , we say \mathcal{D}_1 *stochastically dominates* \mathcal{D}_2 , denoted by $\mathcal{D}_1 \succeq \mathcal{D}_2$, if $\forall a \geq 0$, $\Pr_{X \sim \mathcal{D}_1}(X_1 \geq a) \geq \Pr_{X \sim \mathcal{D}_2}(X_1 \geq a)$. We note an important property of concave functions in the following lemma.

Lemma 2.1 *Given any non-decreasing concave function \mathbf{U} , and three independent non-negative random variables X, Y_1, Y_2 , let \mathcal{D}_1 and \mathcal{D}_2 be the distributions of Y_1 and Y_2 respectively. If $\mathcal{D}_1 \succeq \mathcal{D}_2$, then $\mathbf{E}_{X, Y_1 \sim \mathcal{D}_1} [\mathbf{U}(X + Y) - \mathbf{U}(Y)] \leq \mathbf{E}_{X, Y_1 \sim \mathcal{D}_2} [\mathbf{U}(X + Y) - \mathbf{U}(Y)]$.*

3 Risk Averse Seller: DSIC Mechanisms for Multi-Unit Auctions

In this section, we construct DSIC mechanisms for a risk averse seller. The following theorem summarizes our result.

Theorem 1 *For multi-unit auctions, there is a poly-time computable deterministic SPM with expected utility at least $(1 - \frac{1}{e} - \epsilon)\text{OPT}$, for any $\epsilon > 0$, where OPT is the expected utility of an optimal DSIC mechanism.*

We first prove the existence of an SPM that achieves a $(1 - 1/e)$ -approximation to the optimal expected utility (Section 3.1), however this result is not constructible and does not lead to an efficient implementation. We then identify a set of sufficient properties of $(1 - 1/e - \epsilon)$ -approximate mechanisms that enables us to construct a poly-time algorithm (Section 3.2).

3.1 Existence of a $(1 - 1/e)$ -approximate SPM

Given a set $S = \{x_1, x_2 \dots x_n\}$ of non-negative real number, let $\max_i \{x_1, x_2 \dots x_n\}$ denote the i^{th} largest value in the set, and let it be zero if $i > n$. Let $\mathbf{U}_k : \mathbf{R}^n \rightarrow \mathbf{R}$ be the function defined as $\mathbf{U}_k(S) = \mathbf{U}_k(x_1, x_2 \dots x_n) = \mathbf{U}(\sum_{i=1}^k \max_i \{x_1, x_2 \dots x_n\})$, i.e. utility of the sum of the k largest arguments. Let $\mathbf{U}(S)$ denote $\mathbf{U}_\infty(S) = \mathbf{U}_{|S|}(S)$, the utility of the sum of all variables. We note an important property of \mathbf{U}_k in the following lemma; its proof is deferred to the appendix.

Lemma 3.1 *For any concave utility function \mathbf{U} , and any k and n , the function $\mathbf{U}_k : \mathbf{R}^n \rightarrow \mathbf{R}$ is a symmetric, monotone and submodular.*

We shall use the following *correlation gap* bound established by Agrawal et. al. [1] for monotone submodular functions.

Lemma 3.2 [1] *Given n non-negative random variables X_1, X_2, \dots, X_n with distributions D_1, D_2, \dots, D_n , let D be an arbitrary joint distribution over these n random variables such that the marginal distribution for each X_i remains unchanged. Let D_{ind} be the joint distribution where each X_i is sampled from D_i independent of X_{-i} . Then for any monotone submodular function f over X_1, X_2, \dots, X_n , we have $\frac{\mathbf{E}_{X \sim D_{\text{ind}}} f(X)}{\mathbf{E}_{X \sim D} f(X)} \geq 1 - 1/e$.*

Let \mathcal{M}_{OPT} be a utility optimal DSIC mechanism for a k -unit auction. It follows that in \mathcal{M}_{OPT} , every buyer i is offered a (random) price P_i as a function of other buyers' bids; he receives an item and pays the offered price if and only if his value exceeds the price. The following lemma uses the correlation gap to establish the existence of an SPM which is a $(1 - 1/e)$ -approximation to \mathcal{M}_{OPT} .

Lemma 3.3 *Suppose that \mathcal{M}_{OPT} offers a (random) price P_i to each buyer i (the prices P_i , $1 \leq i \leq n$, may be correlated). Let \mathcal{M}' be a randomized SPM that selects an independent random price P'_i for each buyer, such that P'_i and P_i have the same marginal distribution, and offers items to buyers in decreasing order of prices, until the items run out. Then $\mathbf{U}(\mathcal{M}') \geq (1 - 1/e)\text{OPT}$.*

Proof: Let R_i be the payment obtained in \mathcal{M}_{OPT} from buyer i . Note that P_i and R_i are correlated random variables that depend on the realization of the valuations, and $R_i = P_i$ if $v_i > P_i$, else $R_i = 0$. As at most k buyers can make a positive payment in any realization of \mathcal{M}_{OPT} , we have

$$\mathbf{U}(\mathcal{M}_{\text{OPT}}) = \mathbf{E}[\mathbf{U}(R_1, R_2 \dots R_n)] = \mathbf{E}[\mathbf{U}_k(R_1, R_2 \dots R_n)]$$

Let $R'_i = P'_i$ if $v_i > P'_i$, else $R'_i = 0$. Since the SPM \mathcal{M}' orders buyers in decreasing order of offer prices, so it collects the k largest acceptable prices as payment. We have $\mathbf{U}(\mathcal{M}') = \mathbf{E}[\mathbf{U}_k(R'_1, R'_2 \dots R'_n)]$. Note that R_i and R'_i have the same distribution for each i , except that $R_1, R_2 \dots R_n$ are correlated variables, while $R'_1, R'_2 \dots R'_n$ are mutually independent. Using the submodularity of \mathbf{U}_k (Lemma 3.1) and the correlation gap (Lemma 3.2), we get

$$\mathbf{U}(\mathcal{M}') = \mathbf{E}[\mathbf{U}_k(R'_1, R'_2 \dots R'_n)] \geq (1 - 1/e)\mathbf{E}[\mathbf{U}_k(R_1, R_2 \dots R_n)] = (1 - 1/e)\mathbf{U}(\mathcal{M}_{\text{OPT}})$$

This completes the proof. \square

Correlation gap was used by Yan [16] to show the same approximation ratio for an SPM to expected revenue maximization. However, for revenue maximization, it suffices for the SPM to match a revenue-optimal mechanism only in the probability of sale to each buyer, which solely determines the buyer's contribution to expected revenue. In contrast, for the utility maximization result of Lemma 3.3, *the SPM should match a utility-optimal mechanism in the entire distribution of prices to each buyer*. Also, the SPM for revenue maximization is poly-time computable, since a revenue-optimal mechanism is known (Myerson's mechanism). To the best of our knowledge, the SPM designed in Lemma 3.3 is not poly-time computable: constructing it would need an oracle access to a utility-optimal mechanism. Further, as we have to match \mathcal{M}_{OPT} for each buyer-price pair, guessing the entire price distribution would require time exponential in the number of buyers.

3.2 Algorithm to compute a $(1 - 1/e - \epsilon)$ -approximate SPM

We now present a polynomial time algorithm to compute an SPM whose approximation guarantee essentially matches the existential result above. To simplify the exposition of our algorithm, we assume that prices offered by any truthful mechanism belong to some known set $\mathcal{P} = \{p_1, p_2, \dots\}$ whose size is polynomial in n . Let π_{ij} be the probability that buyer i is offered price p_j in \mathcal{M}_{OPT} . We divide the prices in \mathcal{P} into 3 classes, *small*, *large* and *huge*. Fix some $1 > \epsilon > 0$. Let \mathcal{P}_{hg} be the set of *huge prices* defined as $p_j \geq \mathbf{U}^{-1}(\text{OPT}/\epsilon)$. The distinction between small and large prices depend more intricately on the optimal mechanism. Let p^* be the largest price such that $\sum_{\mathbf{U}^{-1}(\text{OPT}/\epsilon) > p_j \geq p^*} q_j \geq 1/\epsilon^4$, i.e. the threshold where the total sale probability of all large prices add up to at least $1/\epsilon^4$. If such a threshold does not exist, then let $p^* = 0$ (note that p^* must be zero if $k < 1/\epsilon^4$). Let \mathcal{P}_{sm} be all prices less than p^* , so that $\mathcal{P}_{\text{lg}} = \{p_j \mid \mathbf{U}^{-1}(\text{OPT}/\epsilon) > p_j \geq p^*\}$.

In the following lemma, we present a key set of sufficient conditions for a $(1 - 1/e - \epsilon)$ -approximate mechanism which forms the basis of our algorithm; we defer its proof to later in the section.

Lemma 3.4 *Consider any SPM \mathcal{M}' , that offers price p_j to buyer i w.p. π'_{ij} , such that (a) $\sum_{i, p_j \in \mathcal{P}_{\text{sm}}} p_j \pi'_{ij} (1 - F_i(p_j)) = \sum_{p_j \in \mathcal{P}_{\text{sm}}} p_j q_j$, (b) for each $p_j \in \mathcal{P}_{\text{lg}}$, $\sum_i \pi'_{ij} (1 - F_i(p_j)) = q_j$, (c) $\sum_{i, p_j \in \mathcal{P}_{\text{hg}}} \mathbf{U}(p_j) \pi'_{ij} (1 - F_i(p_j)) = \sum_{p_j \in \mathcal{P}_{\text{hg}}} \mathbf{U}(p_j) q_j$ and (d) $\sum_{i, p_j \in \mathcal{P}} \pi'_{ij} (1 - F_i(p_j)) \leq k$. Then we have $\mathbf{U}(\mathcal{M}') \geq (1 - \frac{1}{e} - O(\epsilon))\text{OPT}$.*

Lemma 3.4 states that instead of matching \mathcal{M}_{OPT} in the probability mass of each <buyer, large-price> pair, *it suffices to match the total probability mass at each large price, summed over all buyers*. Thus the *probability mass can be redistributed across buyers* without much loss in utility.

Further, Lemma 3.4 effectively states that the contribution of the small prices and the large prices can be *linearized*. Intuitively, if the small prices make a significant contribution to utility, then the mechanism must be collecting many small prices, so the total revenue from small prices exhibits a concentration around its expectation. Moreover, whenever a huge price is obtained in a realization, we can neglect the contribution from all other buyers in that realization, without losing much of the expected utility. So the contribution of huge prices can be measured separately. *This separation of huge and small prices from large prices enables us to keep the number of distinct large prices to at most a constant*.

Algorithm: We give an outline of the algorithm; the details are provided in the appendix. From Lemma 3.4, it suffices to match \mathcal{M}_{OPT} in (a) the expected revenue from the small prices (\mathcal{R}), (b) the expected contribution to utility from the huge prices H , and (c) the total sale probability at each large price (q_j). The values of these parameters define a *configuration*, and we guess the value of each parameter with appropriate discretization. The number of distinct configurations is bounded by $2^{\text{poly}(1/\epsilon)}$. For each configuration, we check if there exists an SPM satisfying the configuration, using the covering linear program (LP) below. In the LP, the variable x_{ij} denotes the probability that buyer i is offered price p_j .

$$\begin{aligned} \sum_i (1 - F_i(p_j)) x_{ij} &\geq q_j & \forall p_j \in \mathcal{P}_{\text{lg}} \\ \sum_{i, p_j \in \mathcal{P}_{\text{sm}}} (1 - F_i(p_j)) p_j x_{ij} &\geq \mathcal{R} \\ \sum_{i, p_j \in \mathcal{P}_{\text{hg}}} (1 - F_i(p_j)) \mathbf{U}(p_j) x_{ij} &\geq H \\ \sum_{i, j} (1 - F_i(p_j)) x_{ij} &\leq k \\ \sum_j x_{ij} &\leq 1 & \forall j \\ x_{ij} &\in [0, 1] & \forall i, j \end{aligned}$$

Any feasible solution to this linear program gives a distribution of prices for each buyer, which gives us an SPM that satisfies the guessed configuration. We iterate through all the configurations, and pick the best among these SPMs. A deterministic SPM with desired utility guarantees can be easily identified by sampling from this randomized SPM.

3.3 Proof of Lemma 3.4

We begin by introducing two operations on random variables, *split* and *merge*. Using these two operations, we prove two key properties of random variables in Lemmas 3.7 and 3.8; Lemma 3.4 would follow as a corollary of these two lemmas.

Split and Merge Operations: We now define two operations, *merge* and *split*, on non-negative random variables. In the *merge* operation, given a set S of independent non-negative random variables, let X_i, X_j be any two variables in S such that $\mathbf{Pr}[X_i \neq 0] + \mathbf{Pr}[X_j \neq 0] \leq 1$, then variables X_i, X_j are replaced by a new variable Y such that, for each $p > 0$, $\mathbf{Pr}[Y = p] = \mathbf{Pr}[X_i = p] + \mathbf{Pr}[X_j = p]$ and Y is independent of other variables in $S \setminus \{X_i, X_j\}$.

The *split* operation breaks a random variable into a set of independent variables. Formally, given a set S of non-negative (possibly correlated) random variables, first the variables in S are made mutually independent, and then each variable $X_i \in S$ is *split* into an arbitrary pre-specified set of independent random variables $\{X_{i1}, X_{i2}, \dots, X_{it}\}$ such that for each $p > 0$, $\sum_{1 \leq j \leq t} \mathbf{Pr}[X_{ij} = p] =$

$\Pr[X_i = p]$ and the sets of variables created are also made mutually independent. Intuitively, the merge operation introduces negative correlation. Analogously, the split operation introduces independence. In Lemmas 3.5 and 3.6, we establish useful properties of merge and the split operations for a concave non-decreasing function; their proofs are deferred to the appendix.

Lemma 3.5 *Let S be a set of independent non-negative random variables, Let $X_1, X_2 \in S$, and let Y be the variable formed by merging X_1 and X_2 . Then $\mathbf{E}[\mathbf{U}_k(S)] \leq \mathbf{E}[\mathbf{U}_k((S \setminus \{X_1, X_2\}) \cup \{Y\})]$.*

Lemma 3.6 *Consider a sequence of split operations on a set S of arbitrarily correlated non-negative random variables and let S' be the set of independent random variables at the end of the split operation. Then $\mathbf{E}[\mathbf{U}_k(S')] \geq (1 - \frac{1}{e})\mathbf{E}[\mathbf{U}_k(S)]$.*

Using these two operations, we establish an important property in the following lemma, that not only introduces independence across correlated random variables, but also allows to redistribute the probability mass across variables.

Lemma 3.7 *Given an arbitrarily correlated set $S = \{X_1, X_2, \dots, X_n\}$ of non-negative random variables, consider any set $S' = \{X'_1, X'_2, X'_3, \dots, X'_m\}$ of independent non-negative random variables, such that for each value $p_j > 0$, we have $\sum_i \Pr[X_i = p_j] = \sum_i \Pr[X'_i = p_j]$. Then for any concave function \mathbf{U} and any $k > 0$, we have $\mathbf{E}[\mathbf{U}_k(S')] \geq (1 - \frac{1}{e})\mathbf{E}[\mathbf{U}_k(S)]$.*

Proof: We perform the *split* operation on S to create a set $Y = \{Y_{ijl}\}$ of variables as follows: for each $1 \leq i \leq n$ and $p_j > 0$, create $L \rightarrow \infty$ variables $\{Y_{ijl} | 1 \leq l \leq L\}$ where Y_{ijl} takes value p_j w.p. $\frac{\Pr[X_i = p_j]}{L}$ and 0 otherwise. Using Lemma 3.6, we get that $\mathbf{E}[\mathbf{U}_k(Y)] \geq (1 - \frac{1}{e})\mathbf{E}[\mathbf{U}_k(S)]$

Now we perform merge operation repeatedly on variables in Y to simulate variables in S' . The condition in the lemma statement ensures that such merging is always possible, since $L \rightarrow \infty$. Then by Lemma 3.5, we get $\mathbf{E}[\mathbf{U}_k(S')] \geq \mathbf{E}[\mathbf{U}_k(Y)] \geq (1 - 1/e)\mathbf{E}[\mathbf{U}_k(S)]$. \square

The following lemma effectively states that, given a set of independent random variables, the contribution to the utility of huge values can be separated, and for small values, the variables can be replaced by their expectations; we defer its proof to the appendix.

Lemma 3.8 *Given any $\epsilon > 0$ and a set of independent non-negative random variables $S = \{X_1, X_2, X_3 \dots\}$ such that X_i takes value p_i w.p. π_i and 0 otherwise, where $p_1 \geq p_2 \geq p_3 \geq \dots \geq 0$. Also, suppose that $\sum_{X_i \in S} \pi_i \leq k$. Let \hat{p} be a price that satisfies $\hat{p} \geq \mathbf{U}^{-1}(\mathbf{E}[\mathbf{U}_k(S)]/\epsilon)$, and let p^* be any price such that $\sum_{p_i \in [p^*, \hat{p})} \pi_i > \frac{1}{\epsilon^4}$ (p^* is 0, if no such price exists). Also, let $S_{sm} = \{X_i | p_i < p^*\}$, $S_{lg} = \{X_i | p^* \leq p_i < \hat{p}\}$, and $S_{hg} = \{X_i | p_i \geq \hat{p}\}$. Then $\mathbf{E}[\mathbf{U}_k(S)]$ is approximated to within a factor of $(1 \pm O(\epsilon))$ by $\sum_{X_i \in S_{hg}} \mathbf{E}[\mathbf{U}(X_i)] + \mathbf{E}\left[\mathbf{U}\left(\mathbf{E}[S_{sm}] + \sum_{X_i \in S_{lg}} X_i\right)\right]$.*

Now we are ready to prove Lemma 3.4. The revenue from a buyer in a mechanism can be represented by a random variable, possibly correlated with other buyers' random variables. Let \mathcal{M}' be a mechanism that matches \mathcal{M}_{OPT} on the total sale probability for each price, and its sale probability for each <buyer, large-price> pair is same as \mathcal{M} . Using Lemma 3.7, we get $\mathbf{U}(\mathcal{M}') \geq (1 - 1/e)\mathbf{U}(\mathcal{M}_{\text{OPT}})$. As \mathcal{M} and \mathcal{M}' have (approximately) identical revenues from small prices and utilities from huge prices, we can invoke Lemma 3.8 to establish that $\mathbf{U}(\mathcal{M}) \geq (1 - \epsilon)\mathbf{U}(\mathcal{M}')$. This completes the proof.

4 Risk Averse Seller and Risk Averse Buyers: BIC Mechanisms

We design mechanisms when buyers as well as the seller are risk averse, and the seller's objective is to maximize his expected utility. We make two assumptions: (a) for any $i, t < 0$, $\mathcal{U}_i(0) = 0$ and $\mathcal{U}_i(t) = -\infty$; this implies that the mechanism is *ex post individually rational*, and the payment from a buyer is 0 whenever he does not get the item, (b) we further restrict to the set of mechanisms in which payments are always non-negative, i.e. there is no *positive transfer* from the seller to a buyer. Consider any mechanism \mathcal{M} that satisfies these assumptions, let $g(\cdot, \cdot, \cdot)$ and $h(\cdot, \cdot, \cdot)$ be its payment functions. Then we have $g(i, j, v) = 0$ for each i, v and payment p_j , and $h(i, j, v) = 0$ for each i, v and payment $p_j < 0$. Thus, the function $g(\cdot, \cdot, \cdot)$ is not required to describe the mechanism. The following theorem summarizes our result.

Theorem 2 *There exists a polynomial time algorithm to compute a truthful-in-expectation mechanism for a k -unit auction with expected utility at least $(1 - \frac{1}{e})^2 \gamma(k) \text{OPT}$ where OPT is the expected utility of an optimal BIC mechanism. Moreover, for $k \geq 1/\epsilon^3$, there is a $(1 - \frac{1}{e} - \epsilon)$ -approximation.*

Our result giving an improved approximation ratio of $(1 - \frac{1}{e}) \gamma(k)$ for IID buyers is deferred to Appendix B. In the rest of the section, we prove Theorem 2. Let \mathcal{M}_{OPT} be a utility optimal BIC mechanism with these two restrictions, and $h_{\text{OPT}}(\cdot, \cdot, \cdot)$ be its payment function.

Overall Idea: Consider any *soft* randomized sequential mechanism \mathcal{M}' that processes buyers independently (according to its payment function $h'(\cdot, \cdot, \cdot)$), and matches \mathcal{M}_{OPT} for (a) the total probability of each large payment summed over all buyers, (b) the total revenue from small payments and (c) the utility from huge payments. Using techniques developed for DSIC mechanisms, we get $\mathbf{U}(\mathcal{M}') \geq (1 - 1/e) \text{OPT}$. However, converting such soft mechanism into a mechanism that strictly satisfies the inventory constraint while maintaining truthfulness is not easy. In the case of DSIC mechanisms, the buyers were arranged in a decreasing order of prices, noting that *top- k* is a sub-modular function. Here, if we allocate items to buyers with *top- k* payments in a realization of \mathcal{M}' , then the mechanism is no longer truthful. Further, as *first- k* is not a sub-modular function, the desired approximation guarantee cannot be proven if we process buyers according to a fixed order. We get around this problem by constructing a mechanism with $L \rightarrow \infty$ rounds, where in every round, each buyer is processed independently w.p. $1/L$. The revenue from each allocation in this mechanism has an identical distribution. This helps to limit the loss caused by imposing strict inventory constraints. We now describe our mechanism in detail.

The Mechanism: Our mechanism $\mathcal{M}_{\text{rounds}}$ consists of $L \rightarrow \infty$ rounds and $h_{\text{rounds}}(\cdot, \cdot, \cdot)$ is the payment function associated with it. In each round, buyers arrive according to a predefined order. When buyer i arrives, subject to availability of items, he is independently processed with probability $\frac{1}{L}$ as follows: if his reported valuation is v , then he is given an item w.p. $\sum_j h_{\text{rounds}}(i, j, v)$, and whenever he is given an item, he makes a payment of p_j w.p. $\frac{h_{\text{rounds}}(i, j, v)}{\sum_l h_{\text{rounds}}(i, l, v)}$. Once processed, buyer i is not considered for any future rounds. Further, the payment function $h_{\text{rounds}}(\cdot, \cdot, \cdot)$ satisfies following properties:

- (a) $\sum_{i, v, p_j \in \mathcal{P}_{\text{sm}}} p_j h_{\text{rounds}}(i, j, v) f_i(v) = \sum_{i, v, p_j \in \mathcal{P}_{\text{sm}}} p_j h_{\text{OPT}}(i, j, v) f_i(v)$,
- (b) for each $p_j \in \mathcal{P}_{\text{lg}}$, $\sum_{i, v} h_{\text{rounds}}(i, j, v) f_i(v) = \sum_{i, v} h_{\text{OPT}}(i, j, v) f_i(v)$,
- (c) $\sum_{i, v, p_j \in \mathcal{P}_{\text{hg}}} \mathbf{U}(p_j) h_{\text{rounds}}(i, j, v) f_i(v) = \sum_{i, v, p_j \in \mathcal{P}_{\text{hg}}} \mathbf{U}(p_j) h_{\text{OPT}}(i, j, v) f_i(v)$,
- (d) for each i, v, v' , $\sum_j \mathcal{U}_i(v - p_j) h_{\text{rounds}}(i, j, v) \geq \sum_j \mathcal{U}_i(v - p_j) h_{\text{rounds}}(i, j, v')$, and
- (e) $\sum_{i, j, v} h_{\text{rounds}}(i, j, v) f_i(v) \leq k$.

We draw a parallel between the properties of $h_{\text{rounds}}(\cdot, \cdot, \cdot)$ with the algorithm developed in the case of DSIC mechanisms: the first three properties are equivalent to designing a mechanism that matches \mathcal{M}_{OPT} in the total probability for each large payment, the expected revenue from small payments and the expected utility from huge payments. The fourth constraint establishes the truthfulness of $\mathcal{M}_{\text{rounds}}$, and the last constraint ensures its feasibility in expectation. We further note that $\mathcal{M}_{\text{rounds}}$ is *truthful-in-expectation*: conditioned on processing buyer i in some round, the payment function ensures truthfulness in terms of his expected utility.

The following lemma bounds the utility of $\mathcal{M}_{\text{rounds}}$, we defer its proof to later in the section.

Lemma 4.1 *As $L \rightarrow \infty$, $\mathbf{U}(\mathcal{M}_{\text{rounds}}) \geq (1 - \epsilon) \left(1 - \frac{1}{e}\right)^2 \gamma(k) \text{OPT}$.*

Algorithm: To construct an algorithm, we guess the total probability for each large payment (q_j), the utility from huge payments (H) and the revenue from the small payments (\mathcal{R}). The feasibility of a configuration can be checked using a covering LP; the details of the LP are given in the appendix. There are $2^{\text{poly}(1/\epsilon)}$ configurations, and we select a feasible configuration with maximum expected utility. Further, the number of rounds can be limited to $O(n^2)$ with a small loss in the approximation factor. To establish our result, it remains to prove Lemma 4.1.

Proof of Lemma 4.1. Let $\mathcal{I}_{\text{copies}}$ be an instance of the problem where each buyer is split into L independent copies, the copies of buyer i are $i1, i2, \dots, iL$, and the valuation for each copy is drawn independently from F_i . Consider a mechanism $\mathcal{M}_{\text{soft}}$ on $\mathcal{I}_{\text{copies}}$ with L iterations. The l th copy of every buyer is considered in the l th iteration; when buyer il arrives, $\mathcal{M}_{\text{soft}}$ discards him w.p. $(1 - 1/L)$, otherwise it processes him according to $h_{\text{rounds}}(i, \cdot, \cdot)$. In the following lemma, we lower bound the utility of $\mathcal{M}_{\text{soft}}$; its proof follows from Lemma 3.7 and Lemma 3.8.

Lemma 4.2 $\mathbf{U}(\mathcal{M}_{\text{soft}}) \geq (1 - 1/e - \epsilon) \text{OPT}$.

To simplify notation, in the rest of the proof, we refer to the payment function of $\mathcal{M}_{\text{soft}}$ by $h(\cdot, \cdot, \cdot)$. Further, let $k_{\text{exp}} = \sum_{i,j,v} h_{\text{rounds}}(i, j, v) f_i(v)$; note $k_{\text{exp}} \leq k$.

Observe that mechanisms $\mathcal{M}_{\text{rounds}}$ and $\mathcal{M}_{\text{soft}}$ are equivalent with two exceptions: (a) hard inventory constraint of $\mathcal{M}_{\text{rounds}}$, and (b) $\mathcal{M}_{\text{soft}}$ can process more than one copy of a buyer in a realization. We first address the issue of the inventory constraint. Using the correlation gap, we get that the expected number of allocations in $\mathcal{M}_{\text{soft}}$ after first k allocations is at most k_{exp}/e . This alone is not sufficient to prove the lemma as \mathbf{U} is not linear. We note a crucial property of $\mathcal{M}_{\text{soft}}$ in Lemma 4.3, it establishes that *the revenue from any allocation in $\mathcal{M}_{\text{soft}}$ has an identical distribution*. Let \mathcal{D}_i be the distribution on the revenue from first i allocations in $\mathcal{M}_{\text{soft}}$.

Lemma 4.3 *As $L \rightarrow \infty$, we have $\mathbb{P}_{X_i \sim \mathcal{D}_i, X_{i-1} \sim \mathcal{D}_{i-1}}[(X_i - X_{i-1}) = p_j] = \frac{\sum_{i,v} h(i,j,v) f_i(v)}{k_{\text{exp}}}$.*

Proof: Let Y_l be a random variable indicating the revenue made in round l . Clearly Y_l and $Y_{l'}$ have the same distribution for any $l, l' \leq L$. Furthermore, as $L \rightarrow \infty$, conditioned on one allocation in an iteration, the probability of an additional allocation in the same iteration is (almost) 0. Thus we get $\Pr[Y_l = p_j | Y_l \neq 0] = \frac{\sum_{i,v} h(i,j,v) f_i(v)}{k_{\text{exp}}}$. Conditioned on i th allocation happening in round l , the probability that i th allocation has revenue p_j , is exactly $\Pr[Y_l = p_j | Y_l \neq 0]$. This proves the lemma. \square

Consider a new mechanism $\mathcal{M}_{\text{hard}}$ on $\mathcal{I}_{\text{copies}}$ that is identical to $\mathcal{M}_{\text{soft}}$ with an exception that it stops after k allocations. We now bound its utility.

Lemma 4.4 As $L \rightarrow \infty$, $\mathbf{U}(\mathcal{M}_{\text{hard}}) \geq \gamma(k)\mathbf{U}(\mathcal{M}_{\text{soft}})$.

Proof: As payments are non-negative, using Lemma 4.3, we get $\mathcal{D}_i \succeq \mathcal{D}_j$ for each i and $j < i$. The contribution to the utility from the i th allocation in $\mathcal{M}_{\text{soft}}$ is

$$\mathbf{E}_{X \sim \mathcal{D}_i, Y \sim \mathcal{D}_{i-1}} [\mathbf{U}(X) - \mathbf{U}(Y)] = \mathbf{E}_{X \sim \mathcal{D}_1, Y \sim \mathcal{D}_{i-1}} [\mathbf{U}(X + Y) - \mathbf{U}(Y)]$$

Let Z_i denote the above quantity. Using stochastic dominance of \mathcal{D}_i over \mathcal{D}_j for every $j < i$ and Lemma 4.3, we get that $Z_i \leq Z_j$ for any $j < i$.

As the allocations in $\mathcal{M}_{\text{soft}}$ are independent, and the expected number of allocations in $\mathcal{M}_{\text{soft}}$ after first k allocations can be bounded by $k_{\text{exp}}(1 - \gamma(k))$. Let r_1, r_2, \dots, r_n be the probabilities of 1st, 2nd, ..., n th allocation in $\mathcal{M}_{\text{soft}}$. We have $r_i \geq r_{>i}$, thus we get

$$\mathbf{U}(\mathcal{M}_{\text{hard}}) = \sum_{1 \leq i \leq k} r_i Z_i \geq \gamma(k) \left(\sum_{1 \leq i \leq k} r_i Z_i + \sum_{k+1 \leq i \leq n} r_i Z_k \right) \geq \gamma(k) \left(\sum_{1 \leq i \leq n} r_i Z_i \right) \geq \gamma(k) \mathbf{U}(\mathcal{M}_{\text{soft}})$$

This proves the lemma. \square

To bound the utility of $\mathcal{M}_{\text{rounds}}$, we need to address one more issue: $\mathcal{M}_{\text{hard}}$ can process more than one copy of a buyer. Let \mathcal{D}_{i1} be the distribution on the revenue from all copies of first i buyers in $\mathcal{M}_{\text{soft}}$. Let \mathcal{D}_{i2} be the distribution on the revenue from first i buyers in $\mathcal{M}_{\text{rounds}}$. As payments are non-negative, we have $\mathcal{D}_{i1} \succeq \mathcal{D}_{i2}$. Furthermore, for any fixed i , for each l , the distribution on the revenue from the l th allocation among buyer i 's copies is same. The expected number of copies of buyer i processed in $\mathcal{M}_{\text{soft}}$ is 1. Using correlation gap, the expected number of rounds in which buyer i is processed in $\mathcal{M}_{\text{hard}}$ after first processing is $1/e$. Using stochastic dominance of \mathcal{D}_{i1} over \mathcal{D}_{i2} , the expected loss in the utility can be bounded by a factor $1/e$. This completes the proof. \square

Now we give an improved result when $k \geq 1/\epsilon^3$. In a soft randomized sequential mechanism with payment function same as $\mathcal{M}_{\text{rounds}}$, if we discard each buyer independently w.p. ϵ , then w.p. at least $(1 - \epsilon)$, we do not run out of items. The $(1 - 1/e - O(\epsilon))$ -approximation follows by the following property of concave functions.

Lemma 4.5 Let X be a random variable that takes value between 0 and \mathcal{R} for some $\mathcal{R} > 0$. Then, $\mathbf{E}[\mathbf{U}(\mathcal{R} - X)] \geq \left(1 - \frac{\mathbf{E}[X]}{\mathcal{R}}\right) (\mathbf{U}(\mathcal{R}) - \mathbf{U}(0))$ for any non-decreasing concave function \mathbf{U} .

References

- [1] S. Agrawal, Y. Ding, A. Saberi, and Y. Ye. Price of correlations in stochastic optimization. *Operations Research*, 2011. Earlier version in SODA 2010.
- [2] S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*, pages 512–521, 2011.
- [3] S. Alaei, H. Fu, N. Haghpanah, J. D. Hartline, and A. Malekian. Bayesian optimal auctions via multi- to single-agent reduction. *ACM-EC*, 2012.
- [4] K. J. Arrow. The theory of risk aversion. *Aspects of the Theory of Risk Bearing*, 1965.
- [5] T. Chakraborty, E. Even-Dar, S. Guha, Y. Mansour, and S. Muthukrishnan. Approximation schemes for sequential posted pricing in multi-unit auctions. In *WINE*, 2010.

- [6] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010.
- [7] P. Eso and G. Futo. Auction design with risk averse seller. *Economic Letters*, 1999.
- [8] S. Hon-Snir. Utility equivalence in auctions. *The B. E. Journal of Theoretical Economics*, 5(1), 2005.
- [9] E. S. Maskin and J. G. Riley. Optimal auctions with risk averse buyers. *Econometrica*, 1984.
- [10] S. Matthews. Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory*, 30:370–400, 1983.
- [11] S. Matthews. Comparing auctions for risk averse buyers: A buyer’s point of view. *Econometrica*, 55(3):633–646, 1987.
- [12] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [13] J. W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1-2):122136, 1964.
- [14] M. Sundararajan and Q. Yan. Robust mechanisms for risk-averse sellers. In *ACM Conference on Electronic Commerce*, pages 139–148, 2010.
- [15] R. B. Wilson. *Nonlinear Pricing*. Oxford University Press, 1997.
- [16] Q. Yan. Mechanism design via correlation gap. In *SODA*, 2011.

A Omitted Proofs from Section 3

Proof of Lemma 3.1. It is obvious that \mathbf{U}_k is monotone and symmetric. Let $x = (x_1, x_2 \dots x_n)$ and $y = (y_1, y_2 \dots y_n)$, and let $\max\{x, y\} = (z_1, z_2 \dots z_n)$ and $\min\{x, y\} = (z'_1, z'_2 \dots z'_n)$, such that $\forall 1 \leq i \leq n$, $z_i = x_i, z'_i = y_i$ if $x_i \geq y_i$, and $z_i = y_i, z'_i = x_i$ if $x_i < y_i$. To show submodularity we need to show that for any x, y

$$\mathbf{U}_k(\max\{x, y\}) + \mathbf{U}_k(\min\{x, y\}) \leq \mathbf{U}_k(x) + \mathbf{U}_k(y)$$

Or equivalently, that for any $x \geq y$ (i.e. $x_i \geq y_i \forall i$), $1 \leq i \leq n$, and $t \geq x_i \geq y_i$, if $x \wedge_i t = x_1 \dots x_{i-1}, t, x_{i+1} \dots$ and $y \wedge_i t = y_1 \dots y_{i-1}, t, y_{i+1} \dots$, then

$$\mathbf{U}_k(x \wedge_i t) - \mathbf{U}_k(x) \leq \mathbf{U}_k(y \wedge_i t) - \mathbf{U}_k(y) .$$

We shall prove the latter statement. Let $I_k(x) = \sum_{j=1}^k \max_j(x)$. It is easy to see that $I_k(x \wedge_i t) - I_k(x) = \max\{0, t - \max\{x_i, \max_k(x)\}\} \leq \max\{0, t - \max\{y_i, \max_k(y)\}\} = I_k(y \wedge_i t) - I_k(y)$, i.e. I_k itself is submodular. Also, $I_k(x) \geq I_k(y)$. The statement now follows from noting that $\mathbf{U}_k(x) = \mathbf{U}(I_k(x))$, and that $\mathbf{U}(a + b) - \mathbf{U}(b) \leq \mathbf{U}(a' + b') - \mathbf{U}(b') \forall a \leq a', b \geq b'$, for any concave function \mathbf{U} . \square

Proof of Lemma 3.5. The lemma follows from the fact that \mathbf{U}_k is submodular – in fact, it applies to any submodular function. Let $\{p_1, p_2, \dots, p_\ell\}$ be the set of different non-zero values realized by either X_1 or X_2 with non-zero probability. Let π_i, π'_i be the probabilities that X_1 and X_2 realize to p_i , respectively. Fix a realization of random variables in $S \setminus \{X_1, X_2\}$, and let this realization be denoted by \hat{z} ; in this case, we shall express $\mathbf{U}_k(S)$ as $\mathbf{U}_k(X_1, X_2, \hat{z})$ and $\mathbf{U}_k((S \setminus \{X_1, X_2\}) \cup \{Y\})$ as $\mathbf{U}_k(Y, \hat{z})$.

$$\begin{aligned} & \mathbf{E}[\mathbf{U}_k(Y, \hat{z})] - \mathbf{E}[\mathbf{U}_k(X_1, X_2, \hat{z})] \\ &= \sum_i (\pi_i + \pi'_i) \mathbf{U}_k(p_i) - \sum_i \left(\pi_i (1 - \sum_j \pi'_j) + \pi'_i (1 - \sum_j \pi_j) \right) \mathbf{U}_k(p_i) - \sum_{i,j} \pi_i \pi'_j \mathbf{U}_k(p_i + p_j) \\ &= \sum_i \left(\pi_i (\sum_j \pi'_j) + \pi'_i (\sum_j \pi_j) \right) \mathbf{U}_k(p_i) - \sum_{i,j} \pi_i \pi'_j \mathbf{U}_k(p_i + p_j) \\ &= \sum_{i,j} \pi_i \pi'_j (\mathbf{U}_k(p_i) + \mathbf{U}_k(p_j) - \mathbf{U}_k(p_i + p_j)) \geq 0 \end{aligned}$$

The last inequality follows from the definition of submodularity. Since the above inequality holds for any \hat{z} , the lemma follows. \square

Proof of Lemma 3.6. This lemma holds for any monotone submodular function, including \mathbf{U}_k . If a variable $X_i \in S$ is split into $\{X_{i1}, X_{i2}, \dots, X_{it}\}$, then the realization of X_i can be simulated by $\sum_{1 \leq j \leq t} X_{ij}$ with appropriate correlation among them. Thus with appropriate correlation among the variables in S' , we can ensure that $\mathbf{E}[\mathbf{U}_k(S')] = \mathbf{E}[\mathbf{U}_k(S)]$. When variables in S' are made independent, its expected utility is at least $(1 - \frac{1}{e})\mathbf{E}[\mathbf{U}_k(S)]$, by the correlation gap bound in Lemma 3.2, since \mathbf{U}_k is monotone and submodular. \square

Proof of Lemma 3.8. Our proof has 3 steps. Our first step is to separate the contribution of S_{hg} from the rest.

Claim 1 $\mathbf{E}[\mathbf{U}_k(S)]$ is approximated within a factor of $(1 \pm O(\epsilon))$ by $\mathbf{E}[\mathbf{U}_k(S_{\text{sm}} \cup S_{\text{lg}})] + \sum_{X_i \in S_{\text{hg}}} \mathbf{E}[\mathbf{U}(X_i)]$.

Proof: It is easy to upper bound the expected utility: $\mathbf{E}[\mathbf{U}_k(S)] \leq \mathbf{E}[\mathbf{U}_k(S_{\text{sm}} \cup S_{\text{lg}})] + \mathbf{E}[\mathbf{U}_k(S_{\text{hg}})] \leq \mathbf{E}[\mathbf{U}_k(S_{\text{sm}} \cup S_{\text{lg}})] + \sum_{X_i \in S_{\text{hg}}} \mathbf{E}[\mathbf{U}(X_i)]$, since \mathbf{U} is concave.

For the lower bound, note that $\Pr[\mathbf{U}_k(S_{\text{hg}}) > 0] \leq \epsilon$, by Markov's inequality, since $\mathbf{U}_k(S_{\text{hg}}) > 0$ implies $\mathbf{U}_k(S) \geq \mathbf{E}[\mathbf{U}_k(S)]/\epsilon$. Further, the probability that some particular $X_i \in S_{\text{hg}}$ is non-zero, while all other variables in S_{hg} are zero, is at least $\pi_i(1 - \epsilon)$. So we have

$$\begin{aligned} \mathbf{E}[\mathbf{U}_k(S)] &= \mathbf{E}[\mathbf{U}_k(S) - \mathbf{U}_k(S_{\text{hg}})] + \mathbf{E}[\mathbf{U}_k(S_{\text{hg}})] \\ &\geq \Pr[\mathbf{U}_k(S_{\text{hg}}) = 0] \mathbf{E}[\mathbf{U}_k(S_{\text{sm}} \cup S_{\text{lg}}) \mid \mathbf{U}_k(S_{\text{hg}}) = 0] + \mathbf{E}[\mathbf{U}_k(S_{\text{hg}})] \\ &\geq (1 - \epsilon) \mathbf{E}[\mathbf{U}_k(S_{\text{sm}} \cup S_{\text{lg}})] + \sum_{X_i \in S_{\text{hg}}} \pi_i(1 - \epsilon) \mathbf{U}(p_i) \\ &= (1 - \epsilon) \mathbf{E}[\mathbf{U}_k(S_{\text{sm}} \cup S_{\text{lg}})] + (1 - \epsilon) \sum_{X_i \in S_{\text{hg}}} \mathbf{E}[\mathbf{U}(X_i)] \end{aligned}$$

□

Without loss of generality, we may assume that $1/\epsilon$ is an integer. Our second step establishes that for $k > 1/\epsilon^4$, one may simply assume that $k = \infty$.

Claim 2 Let $S' = S_{\text{sm}} \cup S_{\text{lg}}$. If $k > 1/\epsilon^4$, then $(1 + \epsilon) \mathbf{E}[\mathbf{U}_k(S')] \geq \mathbf{E}[\mathbf{U}_{k(1+\epsilon)}(S')] \geq (1 - \epsilon) \mathbf{E}[\mathbf{U}(S')]$.

Proof: The first inequality trivially follows from the definition of \mathbf{U}_k , since $(1 + \epsilon) \mathbf{U}_k(S') \geq \mathbf{U}_{k(1+\epsilon)}(S')$ on every realization.

Since $k > 1/\epsilon^4$ and $\sum_{X_i \in S'} \pi_i \leq k$, so the probability of the event that the number of non-zero variables in a realization exceeds $k(1 + \epsilon)$, is at most ϵ . This follows directly from Chebyshev's inequality. Conditioned on this event, the expected contribution to $\mathbf{U}(S')$ from variables that are positive but does not contribute to $\mathbf{U}_{k(1+\epsilon)}(S')$, is at most $\mathbf{E}[\mathbf{U}(S')]$. This is because the realizations of these *excess positive variables* is independent of the said event. So the difference between $\mathbf{E}[\mathbf{U}(S')]$ and $\mathbf{E}[\mathbf{U}_{k(1+\epsilon)}(S')]$ is at most $\epsilon \mathbf{E}[\mathbf{U}(S')]$, hence the second inequality. □

Our third claim, combined with the first claim, completes the proof of Lemma 3.8.

Claim 3 Let $S' = S_{\text{sm}} \cup S_{\text{lg}}$. Then $\mathbf{E}[\mathbf{U}_k(S')]$ is approximated within a factor of $(1 \pm O(\epsilon))$ by $\mathbf{E}\left[\mathbf{U}\left(\mathbf{E}[S_{\text{sm}}] + \sum_{X_i \in S_{\text{lg}}} X_i\right)\right]$.

Proof: Also, if $k \leq 1/\epsilon^4$, then $S_{\text{sm}} = \emptyset$ by definition, so the claim is trivially true; so we assume that $k > 1/\epsilon^4$. Then, by our second claim above, $\mathbf{E}[\mathbf{U}_k(S')]$ is approximated within a factor of $(1 + 2\epsilon)$ by $\mathbf{E}[\mathbf{U}(S')]$. So it suffices to show that $\mathbf{E}[\mathbf{U}(S')]$ is approximated within a factor of $(1 \pm O(\epsilon))$ by $\mathbf{E}\left[\mathbf{U}\left(\mathbf{E}[S_{\text{sm}}] + \sum_{X_i \in S_{\text{lg}}} X_i\right)\right]$. Again, the upper bound is easy to obtain: $\mathbf{E}[\mathbf{U}(S')] \leq \mathbf{E}\left[\mathbf{U}\left(\mathbf{E}[S_{\text{sm}}] + \sum_{X_i \in S_{\text{lg}}} X_i\right)\right]$ by concavity of \mathbf{U} .

For the lower bound, let us first consider the case that $\mathbf{E}[S_{\text{sm}}] \geq p^*/\epsilon^3$. Then $\mathbf{Var}[S_{\text{sm}}] \leq \epsilon^3 \mathbf{E}[S_{\text{sm}}]$, since $X_i < p^*$ for all $X_i \in S_{\text{sm}}$. Then by Chebyshev's inequality, we have $\sum_{X_i \in S_{\text{sm}}} X_i > (1 - \epsilon) \mathbf{E}[S_{\text{sm}}]$ with probability at least $1 - \epsilon$. So

$$\mathbf{E}[\mathbf{U}(S')] \geq (1 - \epsilon) \mathbf{E}\left[\mathbf{U}\left((1 - \epsilon) \mathbf{E}[S_{\text{sm}}] + \sum_{X_i \in S_{\text{lg}}} X_i\right)\right] \geq (1 - 2\epsilon) \mathbf{E}\left[\mathbf{U}\left(\mathbf{E}[S_{\text{sm}}] + \sum_{X_i \in S_{\text{lg}}} X_i\right)\right].$$

Now let us assume that $\mathbf{E}[S_{\text{sm}}] < p^*/\epsilon^3$. In this case, we show below that

$$\mathbf{E}[\mathbf{U}(S')] \leq \mathbf{E}\left[\mathbf{U}\left(\mathbf{E}[S_{\text{sm}}] + \sum_{X_i \in S_{\text{lg}}} X_i\right)\right] \leq (1 + O(\epsilon))\mathbf{E}[\mathbf{U}(S_{\text{lg}})] \leq (1 + O(\epsilon))\mathbf{E}[\mathbf{U}(S')] .$$

Note that with probability at least $(1 - \epsilon)$, the number of non-zero variables in S_{lg} is at least $(1 - \epsilon)/\epsilon^4$ (by Chebyshev's inequality on the indicator variables of X_i), which implies that $\sum_{X_i \in S_{\text{lg}}} X_i \geq (1 - \epsilon)p^*/\epsilon^4$ – let this event be denoted by A , so that $\Pr[A] \geq (1 - \epsilon)$, and let \bar{A} denote the complementary event. Let $\mathbf{E}[S_{\text{sm}}] = \mathcal{R}_{\text{sm}}$.

$$\begin{aligned} & \mathbf{E}\left[\mathbf{U}\left(\mathcal{R}_{\text{sm}} + \sum_{X_i \in S_{\text{lg}}} X_i\right) - \mathbf{U}(S_{\text{lg}})\right] \\ &= \mathbf{E}\left[\mathbf{U}\left(\mathcal{R}_{\text{sm}} + \sum_{X_i \in S_{\text{lg}}} X_i\right) - \mathbf{U}(S_{\text{lg}}) \mid A\right] \Pr[A] + \mathbf{E}\left[\mathbf{U}\left(\mathcal{R}_{\text{sm}} + \sum_{X_i \in S_{\text{lg}}} X_i\right) - \mathbf{U}(S_{\text{lg}}) \mid \bar{A}\right] (1 - \Pr[A]) \\ &\leq (1 - \epsilon) \left(\mathbf{U}\left(\mathcal{R}_{\text{sm}} + \frac{(1 - \epsilon)p^*}{\epsilon^4}\right) - \mathbf{U}\left(\frac{(1 - \epsilon)p^*}{\epsilon^4}\right)\right) + \epsilon (\mathbf{U}(\mathcal{R}_{\text{sm}} + 0) - \mathbf{U}(0)) \\ &\leq (1 - \epsilon) \frac{\epsilon^4 \mathcal{R}_{\text{sm}}}{(1 - \epsilon)p^*} \mathbf{U}\left(\frac{(1 - \epsilon)p^*}{\epsilon^4}\right) + \epsilon \mathbf{U}(\mathcal{R}_{\text{sm}}) \\ &\leq \epsilon \mathbf{U}\left(\frac{(1 - \epsilon)p^*}{\epsilon^4}\right) + \epsilon \mathbf{U}\left(\frac{p^*}{\epsilon^3}\right) \leq 2\epsilon \mathbf{U}\left(\frac{(1 - \epsilon)p^*}{\epsilon^4}\right) \end{aligned}$$

Both the first and second inequalities follow from concavity of \mathbf{U} , while the last inequality holds for $\epsilon \leq \frac{1}{2}$. Finally, observe that $\mathbf{E}[\mathbf{U}(S_{\text{lg}})] \geq (1 - \epsilon)\mathbf{U}\left(\frac{(1 - \epsilon)p^*}{\epsilon^4}\right)$. \square \square

A.1 Detailed Algorithm for the DSIC Mechanism

From Lemma 3.4, it suffices to guess (a) the threshold p^* separating large and small prices, (b) the total probability of sale q_k for each large price p_k , (c) the expected revenue from small prices \mathcal{R} , and (d) the expected contribution to utility, say h , from huge prices. The large prices (and hence the threshold p^* as well) can be discretized so that they come only from the set $\mathbf{U}^{-1}(\epsilon^5 t \text{OPT})$ for any integer $0 \leq t \leq \frac{1}{\epsilon^6}$. This is because the loss in utility due to discretization, in every realization, is at most $\epsilon^5 \text{OPT}$ per sale at a large price. This is immediate if the large price sale is the first sale to occur in the realization, and concavity implies that the loss can only be smaller if some revenue has already been collected in the realization. So the total loss in expected utility due to this discretization is at most $\epsilon^5 \text{OPT}$ times the expected number of large prices sales. Since there are $O(\frac{1}{\epsilon^4})$ large price sales in expectation, the loss is $O(\epsilon \text{OPT})$.

Similarly, h and $\mathbf{U}(\mathcal{R})$ can be guessed up to a multiple of ϵOPT . We shall guess the total sale probability for each large price to the nearest multiple of ϵ^8 . The total error over all price buckets is less than ϵ^2 and so the loss in the utility due to rounding of the probability masses is bounded by ϵOPT . All these guesses define a *configuration*. The number of configurations is bounded by $2^{\text{poly}(1/\epsilon)}$. For each configuration we check if there exists an SPM satisfying the configuration, using

the covering linear program (LP) below. In this, the variable x_{ij} denotes the probability that buyer i is offered price p_j .

$$\begin{aligned}
\sum_i (1 - F_i(p_j)) x_{ij} &\geq q_j & \forall p_j \in \mathcal{P}_{\text{lg}} \\
\sum_{i, p_j \in \mathcal{P}_{\text{sm}}} (1 - F_i(p_j)) p_j x_{ij} &\geq \mathcal{R} \\
\sum_{i, p_j \in \mathcal{P}_{\text{hg}}} (1 - F_i(p_j)) \mathbf{U}(p_j) x_{ij} &\geq H \\
\sum_{i, j} (1 - F_i(p_j)) x_{ij} &\leq k \\
\sum_j x_{ij} &\leq 1 & \forall j \\
x_{ij} &\in [0, 1] & \forall i, j
\end{aligned}$$

Any feasible solution to this linear program gives a distribution of prices for each buyer, which gives us an SPM that satisfies the guessed configuration. For the configuration that matches \mathcal{M}_{OPT} , any feasible solution of the program gives an SPM that is a $(1 - 1/e - \epsilon)$ -approximation. So our algorithm is to iterate through all the configurations, generate at most one randomized SPM for each configuration, and then simulate them to find out, with high confidence, the best among these SPMs.

B BIC Mechanisms for a Risk Averse Seller and Risk Averse IID Buyers

In this section, we consider the problem of designing a BIC mechanism for a risk averse seller when buyers are risk averse, each buyer's valuation is drawn IID from a known distribution F and their utility functions \mathcal{U} are identical. The following theorem summarizes our result in this setting.

Theorem 3 *There exists a polynomial time algorithm to compute a truthful-in-expectation mechanism with expected utility $(1 - \epsilon) (1 - \frac{1}{e}) \gamma(k) \text{OPT}$, where OPT is the expected utility of the optimal BIC mechanism.*

In the rest of the section, we prove Theorem 3. We call a mechanism to be symmetric if, for all $i \neq i', j, v$, we have $h(i, j, v) = h(i', j, v)$. We first establish an important property of symmetric mechanisms in the following lemma.

Lemma B.1 *There exists a utility optimal BIC mechanism that is symmetric for all buyers.*

Proof: If the optimal BIC mechanism \mathcal{M}_{OPT} is not symmetric, then consider a new mechanism \mathcal{M}' that for every set of reported bids, renames buyers randomly, and runs \mathcal{M}_{OPT} on the these bids. Clearly $\mathbf{U}(\mathcal{M}_{\text{OPT}}) = \mathbf{U}(\mathcal{M}')$. To prove the lemma, it suffices to establish that \mathcal{M}' is a BIC mechanism. Fix any buyer i , his utility in \mathcal{M}' for reporting his true valuation v is

$$\frac{1}{n} \left(\sum_{j, l} (\mathcal{U}(v - p_j) h(l, j, v)) \right)$$

and his utility when he reports v' instead of true valuation v is

$$\frac{1}{n} \left(\sum_{j, l} (\mathcal{U}(v - p_j) h(l, j, v')) \right)$$

From the incentive compatibility property of \mathcal{M}_{OPT} , we get

$$\sum_j (\mathcal{U}(v - p_j)h(l, j, v)) \geq \sum_j (\mathcal{U}(v - p_j)h(l, j, v'))$$

for every l . Summing over l proves the lemma. \square

So from here onwards, we can assume that \mathcal{M}_{OPT} is symmetric; in the optimal mechanism, let $h(j, v)$ be the probability that the buyer gets an item for payment p_j when his valuation is v . Thus the revenue of \mathcal{M}_{OPT} can be represented using the sum of identical (and correlated) random variables X_1, X_2, \dots, X_n , where for each i , $\Pr[X_i = p_j] = \sum_v h(j, v)f(v)$. Let X'_1, X'_2, \dots, X'_n be a set of random variables such that the marginal distribution of X_i and X'_i is same for each i . Let $\mathcal{M}_{\text{soft}}$ be a soft randomized sequential mechanism whose allocation and payment functions are same as \mathcal{M}_{OPT} . Using correlation gap, we get

$$\mathbf{U}(\mathcal{M}_{\text{soft}}) = \mathbf{E} \left[\mathbf{U} \left(\sum_i X'_i \right) \right] \geq \left(1 - \frac{1}{e} \right) \text{OPT}$$

Let \mathcal{M} be the mechanism obtained by applying the inventory constraint on $\mathcal{M}_{\text{soft}}$, i.e. it stops when items run out. We note that \mathcal{M} is still a BIC mechanism: even though the probability of allocations and payments for a buyer for valuation v is not same in $\mathcal{M}_{\text{soft}}$ and \mathcal{M} , these values remain same conditioned on reaching a buyer without running out of inventory. We now bound the expected utility of \mathcal{M} in the following lemma, its proof is similar to the proof of Lemma 4.4.

Lemma B.2 $\mathbf{U}(\mathcal{M}) \geq \gamma(k)\mathcal{M}_{\text{soft}}$.

Thus it suffices to guess the values $h(j, v)$ for each j, v to compute a $(1 - 1/e)^2$ -approximate mechanism. However such algorithm will require exponential time. We need the following lemma to get a polynomial time algorithm.

Lemma B.3 *Consider two randomized and symmetric sequential pricing mechanism \mathcal{M}_1 and \mathcal{M}_2 with inventory constraint of k , and let $h_1(v, j)$ and $h_2(v, j)$ be their payment functions, such that (a) $\sum_{v, p_j \in \mathcal{P}_{sm}} p_j h_1(j, v)f(v) = \sum_{v, p_j \in \mathcal{P}_{sm}} p_j h_2(j, v)f(v)$, (b) for each $p_j \in \mathcal{P}_{lg}$, $\sum_v h_1(j, v)f(v) = \sum_v h_2(j, v)f(v)$, and (c) $\sum_{v, p_j \in \mathcal{P}_{hg}} \mathbf{U}(p_j)h_1(j, v)f(v) = \sum_{v, p_j \in \mathcal{P}_{hg}} \mathbf{U}(p_j)h_2(j, v)f(v)$.*

Its proof is similar to the proof of Lemma 3.4. We note an important step in the proof: if two symmetric mechanisms have the same probability for each payment, then the revenue distribution per buyer remains same, and we do not lose a factor of $(1 - 1/e)$ in the approximation ratio during the split-and-merge operation.

Thus per buyer values of probability for each large payment (q_j), the utility from huge payments (H) and the revenue from the small payments (\mathcal{R}); the feasibility of a configuration can be checked using an LP.

C Omitted Details from Section 4

BIC Mechanisms for a Risk Averse Sellers: LP to check the feasibility of a configuration

$$\begin{aligned}
\sum_j \mathcal{U}_i(v - p_j) h_{\text{rounds}}(i, j, v) &\geq \sum_j \mathcal{U}_i(v - p_j) h_{\text{rounds}}(i, j, v') && \forall i, v, v' \\
\sum_{i,v} h_{\text{rounds}}(i, j, v) f_i(v) &\geq q_j && \forall p_j \in \mathcal{P}_{\text{lg}} \\
\sum_{v, p_j \in \mathcal{P}_{\text{sm}}} p_j f_i(v) h_{\text{rounds}}(i, j, v) &\geq \mathcal{R} \\
\sum_{v, p_j \in \mathcal{P}_{\text{hg}}} \mathbf{U}(p_j) f_i(v) h_{\text{rounds}}(i, j, v) &\geq H \\
\sum_{i,j,v} h_{\text{rounds}}(i, j, v) f_i(v) &\leq k \\
\sum_j h_{\text{rounds}}(i, j, v) &= 1 && \forall i, v \\
h_{\text{rounds}}(i, j, v) &\in [0, 1] && \forall i, j, v
\end{aligned}$$

D Mechanisms for a Risk Neutral Seller and Risk Averse Buyers

In this section, we design mechanisms for a risk neutral seller when buyers are risk averse. The following theorem summarizes our result; it shows that (a) optimal BIC mechanism can be designed for risk averse seller and (b) the gap between an optimal truthful-in-expectation and an optimal BIC mechanisms is bounded by $\gamma(k)$.

Theorem 4 *Given risk averse buyers and a risk neutral seller, let OPT be the expected revenue of the optimal BIC mechanism. Then (a) there exists a polynomial time algorithm to compute a truthful-in-expectation mechanism with expected revenue $\gamma(k)\text{OPT}$, and (b) there exists a polynomial time algorithm to compute a BIC mechanism with expected revenue OPT .*

Proof: We begin by proving the first part of the theorem. We note that the objective of the following LP is at least OPT . The first constraint in the LP ensures that the mechanism is truthful, the second constraint ensures that exactly one allocation-payment event realizes in expectation for a buyer's bid, and the third constraint is the inventory constraint.

$$\begin{aligned}
&\text{Maximize } \sum_{i,j,v} (g(i, j, v) p_j f_i(v) + h(i, j, v) p_j f_i(v)) \\
&\sum_j (\mathcal{U}_i(v - p_j) h(i, j, v) + \mathcal{U}_i(-p_j) g(i, j, v)) \geq \sum_j (\mathcal{U}_i(v - p_j) h(i, j, v') + \mathcal{U}_i(-p_j) g(i, j, v')) && \forall i, v, v' \\
&\sum_j (h(i, j, v) + g(i, j, v)) = 1 && \forall i, v \\
&\sum_{i,j,v} h(i, j, v) f_i(v) \leq k \\
&h(i, j, v), g(i, j, v) \in [0, 1] && \forall i, j, v
\end{aligned}$$

We construct a mechanism from the LP solution using Algorithm 1. We note that the mechanism is truthful-in-expectation: the position of buyer i in the sequence is independent of his valuation, and conditioned on reaching him without running out of items, he allocation and payment probabilities are equal to the LP solution. We now establish the approximation bound.

For the purpose of analysis, we replace the payment from buyer i with a deterministic quantity $\text{payment}_i = \sum_{j,v} (g(i, j, v) p_j f_i(v) + h(i, j, v) p_j f_i(v))$. Let $\text{prob}_i = \sum_{j,v} (g(i, j, v) f_i(v) + h(i, j, v) f_i(v))$, i.e. the probability of allocation to buyer i . Further, let $k_{\text{exp}} = \sum_{i,j,v} h(i, j, v) f_i(v)$. W.l.o.g., assume that the mechanism processes buyers in order $1, 2, 3, \dots, n$. Let r_i be the probability that k items are allocated before buyer i arrives. We have $r_1 \leq r_2 \leq \dots \leq r_n$, and

$$\sum_i r_i \times \text{prob}_i \leq (1 - \gamma(k)) \times k$$

Algorithm 1 Truthful in Expectation Mechanism for a Risk Neutral Seller and Risk Averse Buyers

Arrange the buyers in the decreasing order of $\frac{\sum_{j,v} (g(i,j,v)p_j f_i(v) + h(i,j,v)p_j f_i(v))}{\sum_{j,v} (g(i,j,v)f_i(v) + h(i,j,v)f_i(v))}$.

for each buyer i in the sequence **do**

if there is at least one item remaining **then**

 Ask buyer i for his valuation. Let v be his reported valuation.

 Allocated an item to him w.p. $\sum_j h(i,j,v)$, and he makes a payment of p_j w.p. $\frac{h(i,j,v)}{\sum_l h(i,l,v)}$ in this case. When he is not given an item, he makes a payment of p_j w.p. $\frac{g(i,j,v)}{\sum_l g(i,l,v)}$.

end if

end for

Let $t_i = \frac{\text{payment}_i}{\text{prob}_i}$. We get $t_1 \geq t_2 \geq \dots \geq t_n$. The revenue of the mechanism is

$$\begin{aligned} & \sum_i (1 - r_i) \times \text{payment}_i \\ &= \sum_i (1 - r_i) \times t_i \times \text{prob}_i \\ &\geq \gamma(k) \sum_i \text{payment}_i \end{aligned}$$

where the last inequality follows using $t_1 \geq t_2 \geq \dots \geq t_n$, and $\sum_i r_i \times \text{prob}_i \leq (1 - \gamma(k)) \times k$. This completes the proof of the first claim.

Now we consider the second part of the theorem. Its proof is similar to the first part; we consider an LP as earlier, however we also add a separation oracle for the Border's inequality; [3] show that such separation oracle can be designed in polynomial time. The constructive proof of the Border's inequality also gives a way of allocating items to buyers based on their valuations honouring supply constraints. When an item is allocated to buyer i , then he is charged p_j w.p. $\frac{h(i,j,v)}{\sum_j h(i,j,v)}$. Similarly, when he is not allocated an item, he pays p_j w.p. $\frac{g(i,j,v)}{\sum_j g(i,j,v)}$. Thus the payment functions of a buyer can be implemented exactly. \square

E DSIC Mechanisms for Risk Averse Sellers: Unit Demand Buyers

In this section, we consider the problem of designing DSIC mechanisms for a risk averse sellers in a multi-parameter settings where buyers are unit-demand. We are able to combine our techniques with a recent approximation result for revenue maximization for the unit-demand problem, when there are multiple distinct items. We note that for such multi-parameter settings, no characterization of an optimal mechanism is known even for the expected revenue objective. This gives our second result.

Theorem 5 *For unit-demand setting with multiple distinct items, there is a poly-time computable randomized order-independent sequential posted price mechanism with expected utility at least $\left(\frac{1-1/e}{6.75} - \epsilon\right) \text{OPT}$, for any $\epsilon > 0$, where OPT is the expected utility of an optimal DSIC mechanism.*

Our obtain our result by combining our tools for handling risk aversion with a recent result of Chawla et. al. [6], who gave a $\frac{1}{6.75}$ -approximation to this problem for the expected revenue objective. The approximation factor has since been improved to $\frac{1}{4}$ for expected revenue, by Alaei [2]. To the best of our knowledge, our algorithm gives the first non-trivial approximation guarantee for maximizing expected (concave) utility in any multi-parameter setting.

In the rest of the section, we prove Theorem 5. We will reduce the problem of designing a mechanism for any instance \mathcal{I} of a unit-demand setting with n buyers and m items to finding a special type of mechanism for an instance \mathcal{I}' of a single parameter setting with *matrix feasibility constraint*. Matrix feasibility constraint is defined as follows: there are mn buyers arranged in an $n \times m$ matrix, and the seller can serve at most one buyer from each row and at most one buyer from each column. There is a natural mapping between the two instances: split buyer $i \in \mathcal{I}$ into m independent buyers $i1, i2, \dots, im$ and the distribution on the valuation of buyer ij is same as the distribution of buyer i 's valuation for item j . Clearly any feasible assignment of items in \mathcal{I} corresponds to a feasible assignment of items in \mathcal{I}' and vice versa.

We define a class of pricing mechanisms, called *partially order-oblivious posted pricing (partial-OPM)*, for the single parameter setting with matrix constraint: the mechanism specifies an ordering σ on the rows on the matrix, along with independent price distributions for all buyers. All buyers in row $\sigma(i)$ (i^{th} row in the sequence) of the matrix arrive before any buyer in row $\sigma(i+1)$; however, the order of buyers in each row $\sigma(i)$ is chosen adversarially, where the adversary *knows the valuation of buyers and the mechanism's prices and wishes to minimize the utility obtained*. Partial-OPM in \mathcal{I}' can be mapped naturally to an SPM for \mathcal{I} : buyers arrive in the order specified by σ , and if buyer ij was offered a price of x in \mathcal{I}' , then buyer i is offered item j at a price x (if item j is available when buyer i arrives). The following lemma describes the relation between these two types of problems.

Lemma E.1 *Given a unit-demand instance \mathcal{I} and its corresponding single-parameter instance \mathcal{I}' as defined above, if OPT and OPT' are the optimal expected utilities for \mathcal{I} and \mathcal{I}' respectively, then $\text{OPT}' \geq \text{OPT}$. Furthermore, given any partial-OPM for \mathcal{I}' , there is a poly-time computable SPM for \mathcal{I} with expected utility at least as much as the expected utility of the partial-OPM for \mathcal{I}' .*

Proof: The scheme for designing an approximate utility-optimal mechanism mirrors the scheme for designing approximately revenue-optimal mechanism in [6], as follows. Fix any utility optimal deterministic mechanism \mathcal{M}_{OPT} for \mathcal{I} . Note that, any deterministic truthful multi-parameter mechanism for unit demand buyers can be interpreted as offering a price menu with one price for each item/service to each buyer as a function of other buyers' bids [15]. We construct a corresponding mechanism \mathcal{M}' on \mathcal{I}' that has the same allocation function as in [6]. In *every realization of the bid vector*, the revenue in \mathcal{M}' is *at least as large as* the revenue in \mathcal{M}_{OPT} . Thus the utility of the former is at least the utility of the latter. So the optimal expected utility on \mathcal{I}' is at least as much as the optimal expected utility on \mathcal{I} .

Next, similar to [6], we show that we can use a partial-OPM \mathcal{M}' for \mathcal{I}' to construct an item pricing mechanism \mathcal{M} for \mathcal{I} . Order the buyers in \mathcal{I} as per the order of rows of matrix in \mathcal{M}' , and fix a realization of buyers' valuations and prices in \mathcal{M}' . In the corresponding realization in \mathcal{I} , when buyer $i \in \mathcal{I}$ arrives, he is shown the set of available items and the price asked for item j is the price asked for buyer $ij \in \mathcal{I}'$. As the buyer i chooses the item that maximizes his utility, the revenue from buyer i is no less than the revenue from buyers in row i in \mathcal{I}' (since the mechanism in \mathcal{I}' is order-oblivious within a row). Thus the revenue obtained in \mathcal{M} is at least the minimum revenue (among all orderings of buyers within each row) in \mathcal{M}' , for *every realization of prices and*

valuations. Hence the expected utility of \mathcal{M} is at least as much as the expected utility of \mathcal{M}' . This completes the proof of the lemma. \square

It is worth noting that Lemma E.1 deviates from the reduction used by Chawla et. al. [6] from unit-demand setting to a single-parameter setting for expected revenue in the following manner: Chawla et. al. [6] computed a *completely order-oblivious* posted pricing (where even buyers in different rows can be intermingled adversarially) for \mathcal{I}' , which yields a *completely order-oblivious* posted pricing for \mathcal{I} . This implication remains true for utility as well. However, we are not able to analyze the performance of a complete OPM in the case of expected utility; instead, we are able to compute a partial-OPM with a good approximation guarantee on expected utility in \mathcal{I}' . As will be clear below, the only part where we need to restrict ourselves to partial-OPMs is Lemma E.3. Its analogue for revenue holds even for complete OPMs by a simple application of Markov's inequality [6], but the proof is more complicated for utility maximization.

In the rest of the section, we prove the following theorem, which combined with Lemma E.1 implies Theorem 5.

Theorem 6 *For the single parameter setting with matrix feasibility constraint, there is a polynomial time algorithm to compute a partial-OPM with a given ordering among the rows and expected utility $\left(\frac{1-1/e}{6.75} - \epsilon\right)$ OPT. Furthermore, the price distributions computed by the algorithm yields this guarantee for any ordering σ among the rows.*

Let \mathcal{M}_{OPT} be an optimal deterministic truthful mechanism for a single-parameter instance with matrix constraint, so by incentive compatibility, every buyer is offered a price as a function of other buyers' bids. Let $\mathcal{P} = \{p_1, p_2, \dots\}$ be the set of distinct prices offered in any mechanism. Let π_{ijk} be the probability that buyer ij is offered price p_k in \mathcal{M}_{OPT} . We note the important properties of prices in \mathcal{M}_{OPT} : (a) $\sum_{j,k} \pi_{ijk} (1 - F_{ij}(p_k)) \leq 1 \forall i$ (at most one buyer is served from the i^{th} row), (b) $\sum_{i,k} \pi_{ijk} (1 - F_{ij}(p_k)) \leq 1 \forall j$ (at most one buyer is served from the j^{th} column), and (c) $\sum_k \pi_{ijk} \leq 1 \forall i, j$ (buyer ij is offered only one price in every realization).

We define *soft-OPM* as a type of posted pricing mechanism, where (a) the prices for buyers are drawn independently, (b) it allocates items to all buyers who have valuations above their respective offered prices, and (c) satisfies feasibility constraints *in expectation* but not in every realization. In other words, a soft-OPM sets prices so that it sells to at most one buyer in expectation in each column and in each row. It is obvious that the utility obtained by a soft-OPM is independent of the arrival order of buyers, and the revenue from a buyer is independent of past or future events in a realization of the mechanism.

We note an important property of soft-OPMs in the following lemma. It enables the redistribution of probability mass for any given price across buyers, and makes it sufficient to guess the total sale probability at each price. It is the analogue of Lemma 3.7 for multi-unit auctions. Its proof is practically identical to that of Lemma 3.7 – it applies a series of split operations followed by a series of merge operations to transform an optimal mechanism to a given soft-OPM. The only difference is that the submodular function in this case is \mathbf{U}_∞ instead of \mathbf{U}_k , that is, there is no explicit bound on the inventory. So we skip the proof of this lemma.

Lemma E.2 *Consider any soft-OPM \mathcal{M}_1 , that offers price $p_k \in \mathcal{P}$ with probability π'_{ijk} , and satisfies the following condition: $\sum_{i,j} \pi'_{ijk} (1 - F(p_k)) = \sum_{i,j} \pi_{ijk} (1 - F(p_k)) = q_k, \forall p_k \in \mathcal{P}$ (for every price, the total sale probability, i.e. the sale probabilities summed over all buyers, at a given price in \mathcal{M}_1 is the same as that in \mathcal{M}_{OPT}). Then $\mathbf{U}(\mathcal{M}_1) \geq (1 - \frac{1}{e})\text{OPT}$.*

Algorithm: We shall design a polynomial time algorithm to compute a soft-OPM \mathcal{M}_1 such that $\mathbf{U}(\mathcal{M}_1) \geq (1 - \frac{1}{e} - O(\epsilon))\text{OPT}$. We shall then construct a partial-OPM \mathcal{M}_2 from \mathcal{M}_1 , as follows: if buyer ij is offered price p_k w.p. π'_{ijk} in \mathcal{M}_1 , then \mathcal{M}_2 offers price p_k to ij w.p. $\pi'_{ijk}/3$ and a price of ∞ with the remaining probability (provided that serving buyer ij will not violate feasibility of matrix constraint). Lemma E.3 below implies that \mathcal{M}_2 gives us the required guarantee, so it only remains to design an algorithm to compute a soft-OPM. To compare our result with the $1/6.75$ -approximation for expected revenue [6], one may note that for linear utility, a reduction to soft-OPM such as Lemma E.2 holds quite easily without losing a factor of $(1 - \frac{1}{e})$, while Lemma E.3 naturally remains true (moreover, it even holds for complete OPMs).

Lemma E.3 *For any soft-OPM \mathcal{M}_1 , if we construct a partial-OPM \mathcal{M}_2 from \mathcal{M}_1 as above, then $\mathbf{U}(\mathcal{M}_2) \geq \mathbf{U}(\mathcal{M}_1)/6.75$ for any ordering among the rows.*

Proof: We first note that a soft-OPM \mathcal{M}_3 that offers price p_k to buyer ij w.p. $\pi'_{ijk}/3$ has expected utility $\mathbf{U}(\mathcal{M}_3) \geq \mathbf{U}(\mathcal{M}_1)/3$ since \mathbf{U} is concave. We shall now show that $\mathbf{U}(\mathcal{M}_2) \geq \frac{4}{9}\mathbf{U}(\mathcal{M}_3)$.

Given an order σ on the rows of the matrix, consider an arbitrary buyer ij . Fix a realization of prices and valuations (\mathbf{p}, \mathbf{v}) of all buyers that belong to rows $\sigma(k)$ for $k < \sigma^{-1}(i)$ (buyers from previous rows than row i according to σ), other than buyers belonging to column j , let this be event $E_{\mathbf{p}, \mathbf{v}}$. We define event $E'_{\mathbf{p}, \mathbf{v}}$ as the event where $E_{\mathbf{p}, \mathbf{v}}$ occurs, and in addition, and in the set of buyers that arrive before ij , no buyer from the i^{th} row or j^{th} column has enough valuation to accept its offered price. If $E'_{\mathbf{p}, \mathbf{v}}$ occurs, then the seller will definitely be able to sell to buyer ij if the latter can accept the offered price.

We note that in event $E'_{\mathbf{p}, \mathbf{v}}$, the revenue generated by \mathcal{M}_2 from all buyers that come before ij is a fixed value, let it be $\mathcal{R}_{\mathbf{p}, \mathbf{v}}$. This is because the valuations and prices for all buyers that come before ij , other than those in row i and column j , are fixed, and there is no sale in row i and column j until buyer ij arrives. Let $X_{\mathbf{p}, \mathbf{v}}$ be a random variable indicating the revenue generated by \mathcal{M}_3 in event $E_{\mathbf{p}, \mathbf{v}}$ before buyer ij arrives. As \mathcal{M}_3 is a soft-OPM, every sale that happens in \mathcal{M}_2 in event $E'_{\mathbf{p}, \mathbf{v}}$ also happens in \mathcal{M}_3 in event $E_{\mathbf{p}, \mathbf{v}}$. So we get $X_{\mathbf{p}, \mathbf{v}} \geq \mathcal{R}_{\mathbf{p}, \mathbf{v}}$ in every realization.

The probability that no buyer in the i^{th} row, except perhaps ij , has a valuation exceeding its offered price, is at least $2/3$, by Markov's inequality. The same holds for the j^{th} column. Moreover, since these two events are independent, we get that $\Pr[E'_{\mathbf{p}, \mathbf{v}}] \geq \frac{4\Pr[E_{\mathbf{p}, \mathbf{v}}]}{9}$.

Let R_{ij} be a random variable that indicates the revenue from buyer ij in \mathcal{M}_3 . We note that R_{ij} is independent of $E_{\mathbf{p}, \mathbf{v}}$. The revenue from ij in \mathcal{M}_2 has same distribution as R_{ij} in event $E'_{\mathbf{p}, \mathbf{v}}$, and it is 0 otherwise. The contribution to utility of buyer ij in \mathcal{M}_3 is

$\sum_{\mathbf{p}, \mathbf{v}} (\mathbf{U}(R_{ij} + X_{\mathbf{p}, \mathbf{v}}) - \mathbf{U}(X_{\mathbf{p}, \mathbf{v}})) \Pr[E_{\mathbf{p}, \mathbf{v}}]$. The contribution to utility of buyer ij in \mathcal{M}_2 is $\sum_{\mathbf{p}, \mathbf{v}} (\mathbf{U}(R_{ij} + \mathcal{R}_{\mathbf{p}, \mathbf{v}}) - \mathbf{U}(\mathcal{R}_{\mathbf{p}, \mathbf{v}})) \Pr[E'_{\mathbf{p}, \mathbf{v}}]$. Since \mathbf{U} is concave, so in every realization,

$$(\mathbf{U}(R_{ij} + \mathcal{R}_{\mathbf{p}, \mathbf{v}}) - \mathbf{U}(\mathcal{R}_{\mathbf{p}, \mathbf{v}})) \geq (\mathbf{U}(R_{ij} + X_{\mathbf{p}, \mathbf{v}}) - \mathbf{U}(X_{\mathbf{p}, \mathbf{v}}))$$

So buyer ij 's contribution to utility in \mathcal{M}_2 is at least $4/9$ times his contribution in \mathcal{M}_3 . Summing over contributions of all buyers in every event, we get that $\mathbf{U}(\mathcal{M}_2) \geq \frac{4}{9}\mathbf{U}(\mathcal{M}_3)$. \square

Algorithm to compute soft-OPM: It is similar to the algorithm designed for computing SPM for k -unit auction.

We divide the prices in \mathcal{M}_{OPT} into 3 classes, *small*, *large* and *huge*. Let \mathcal{P}_{hg} be the set of *huge* prices $p_k \geq \mathbf{U}^{-1}(\text{OPT}/\epsilon)$. The distinction between small and large prices depend more intricately on the optimal mechanism. Let p^* be the largest price such that $\sum_{\mathbf{U}^{-1}(\text{OPT}/\epsilon) > p_k \geq p^*} q_k \geq 1/\epsilon^4$,

i.e. the threshold where the total sale probability of all large prices add up to at least $1/\epsilon^4$. If such a threshold does not exist, then let $p^* = 0$. Let \mathcal{P}_{sm} be all prices less than p^* , so that $\mathcal{P}_{\text{lg}} = \{p_k \mid \mathbf{U}^{-1}(\text{OPT}/\epsilon) > p_k \geq p^*\}$. Now, using Lemma 3.8 in the proof of Lemma E.2 allows us to prove the following variant of Lemma E.2. The proof of Lemma E.4 is practically identical to that of Lemma 3.4, so we omit it.

Lemma E.4 *Consider any soft-OPM \mathcal{M}_1 , that offers price p_k to buyer ij w.p. π'_{ijk} , such that (a) $\sum_{i,j,p_k \in \mathcal{P}_{\text{sm}}} p_k \pi'_{ijk} (1 - F_{ij}(p_k)) = \sum_{p_k \in \mathcal{P}_{\text{sm}}} p_k q_k$, (b) for each $p_k \in \mathcal{P}_{\text{lg}}$, $\sum_{i,j} \pi'_{ijk} (1 - F_{ij}(p_k)) = q_k$, and (c) $\sum_{i,j,p_k \in \mathcal{P}_{\text{hg}}} \mathbf{U}(p_k) \pi'_{ijk} (1 - F_{ij}(p_k)) = \sum_{p_k \in \mathcal{P}_{\text{hg}}} \mathbf{U}(p_k) q_k$. Then we have $\mathbf{U}(\mathcal{M}_1) \geq (1 - \frac{1}{e} - O(\epsilon))\text{OPT}$.*

Thus it suffices to guess (a) the threshold p^* separating large and small prices, (b) the total probability of sale q_k for each large price p_k , (c) the expected revenue from small prices \mathcal{R} , and (d) the expected contribution to utility, say h , from huge prices. Just as in the case of multi-unit auctions, the large prices (and hence the threshold p^* as well) can be discretized so that they come only from the set $\mathbf{U}^{-1}(\epsilon^5 t \text{OPT})$ for any integer $0 \leq t \leq \frac{1}{\epsilon^6}$. Similarly, h and $\mathbf{U}(\mathcal{R})$ can be guessed up to a multiple of ϵOPT , and the total sale probability for each large price to the nearest multiple of ϵ^8 . All these guesses define a *configuration*. The number of configurations is bounded by $2^{\text{poly}(1/\epsilon)}$.

$$\begin{aligned}
\sum_{i,j} (1 - F_{ij}(p_k)) x_{ijk} &\geq q_k \quad \forall p_k \in \mathcal{P}_{\text{lg}} \\
\sum_{i,j,p_k \in \mathcal{P}_{\text{sm}}} (1 - F_{ij}(p_k)) p_k x_{ijk} &\geq \mathcal{R} \\
\sum_{i,j,p_k \in \mathcal{P}_{\text{hg}}} (1 - F_{ij}(p_k)) \mathbf{U}(p_k) x_{ijk} &\geq h \\
\sum_{i,k} (1 - F_{ij}(p_k)) x_{ijk} &\leq 1 \quad \forall j \\
\sum_{j,k} (1 - F_{ij}(p_k)) x_{ijk} &\leq 1 \quad \forall i \\
\sum_k x_{ijk} &\leq 1 \quad \forall i, j \\
x_{ijk} &\geq 0 \quad \forall i, j, k
\end{aligned}$$

In the above LP, the variable x_{ijk} denotes the probability that buyer ij is shown price p_k . Each feasible configuration gives us a soft-OPM and we select the one with maximum utility.